Supplemental Material

Supplemental Material S1: Preprocessing; Related to Fig. 1

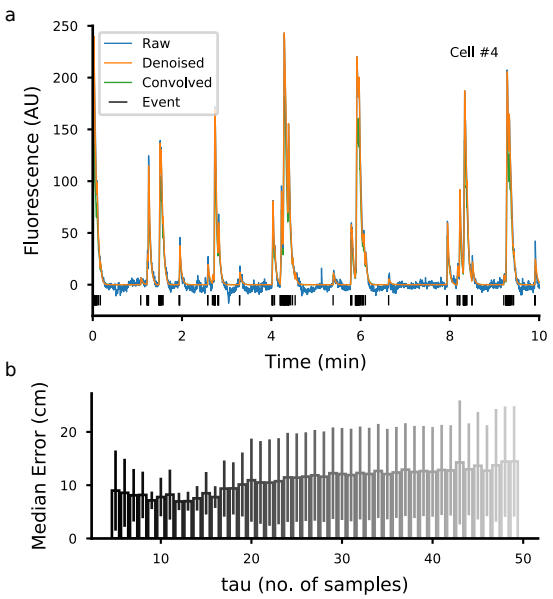

Figure S1: Preprocessing. Here we illustrate the type of signals extracted using CNMF-e that we used for our analysis. The algorithm's output is composed of a spatial footprint, the raw calcium signal, the denoised signal, the calcium events timings for each cell and their magnitude. **a**) Example of fluorescence signal extracted using CNMF-e (Zhou et al., 2018). The raw, denoised and event signals correspond to the output of the algorithm. The raw signal corresponds to the fluorescence levels attributed to the single sources when the contaminating signal due to neuropil in the background and to neighbouring sources have been eliminated. This signal includes noise due to the recording apparatus. The convolved signal is the one used in this work and is obtained using an auto-regressive model that uses the estimated calcium event timings. The parameter of the model as well as the calcium event times, the spatial footprints and the background estimation are all part of the minimization process in CNMF-e (Zhou et al., 2018). In our analysis we used the calcium events times because these gave us the closest estimates of each cell's activity. We then convolved the calcium events in time with a decaying temporal profile to cumulate information in time. Similar procedures have been used in the past to amplify the signal with temporally sparse data (Ziv et al., 2013). **b**) Median Error (mean $\pm$ s.d.) of the position decoding for different time scales of the exponential kernel used to convolve the calcium events. The chosen time scale corresponds to the one that minimizes the position decoding error (2.2 s in this example, 12 samples).

Supplemental Material S2: Place-cells; Related to Fig. 4

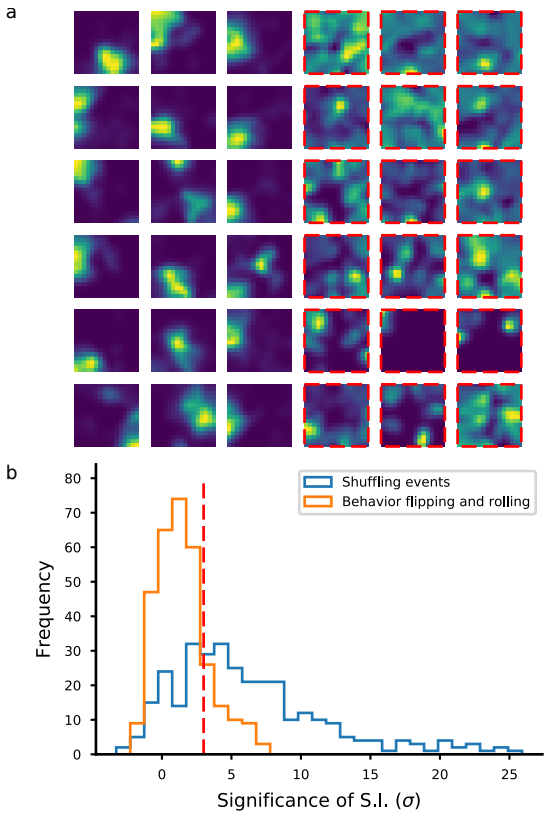

Figure S2: Place-cells. **a**) Gaussian-smoothed event density maps for a selection of cells, normalized by the mouse's occupancy time per 16x16 unit area and the cell's maximum response in the 50 cm x 50 cm arena (Ziv et al., 2013). The cells have been ordered by the statistical significance of their spatial information, from the most to the least significant. Here we show the group of 18 cells with the most significant SSI (place cells) and the 18 with the least significant SSI (non place-cells, red-borders) from the dentate gyrus of a representative mouse (DG3). **b**) Histogram of SSIs for all the DG cells recorded for mouse DG3 computed with the event shuffling method (blue) and the trajectory shuffling method (orange). In order to quantify the spatial tuning of the cells, we applied standard methods used in electrophysiology to measure the spatial information contained in the activity of each cell using the calcium events (Skaggs et al., 1993). We corrected for the sampling bias problem (Panzeri et al., 2007) by using a shuffling approach. We quantified the tuning of the cells by the statistical significance of the spatial information content (SSI). Similar to what we did for assessing the chance level decoding (see Fig. S5), we used two shuffling methods for assessing the SSI of each cell. The first method (blue histogram) consisted in shuffling the event timings, as it is commonly used in the definition of a place cell. We identified place cells as those with an SSI>3. For comparison, here we report a second method which consisted in time reversing the animal's trajectory ("flip") and then shifting it in time by a random amount ("roll"). This second method preserves the activity of all cells while destroying its association to the position of the animal and therefore it's more conservative (see also Fig. S3). The latter method highlights the importance of adopting the correct null hypothesis to assess the spatial information of the cells. In our recordings, many of the cells that are classically classified as place cells have a spatial information content that is not statistically different from the one obtained by a conservative shuffling of the trajectory which still disrupts the association between a cell's activity and the mouse position at each time point. In the text, we

used the former definition for consistency with the literature.

Supplemental Material S3: Fraction of place cells and activity levels in DG and CA1; Related to Fig. 4

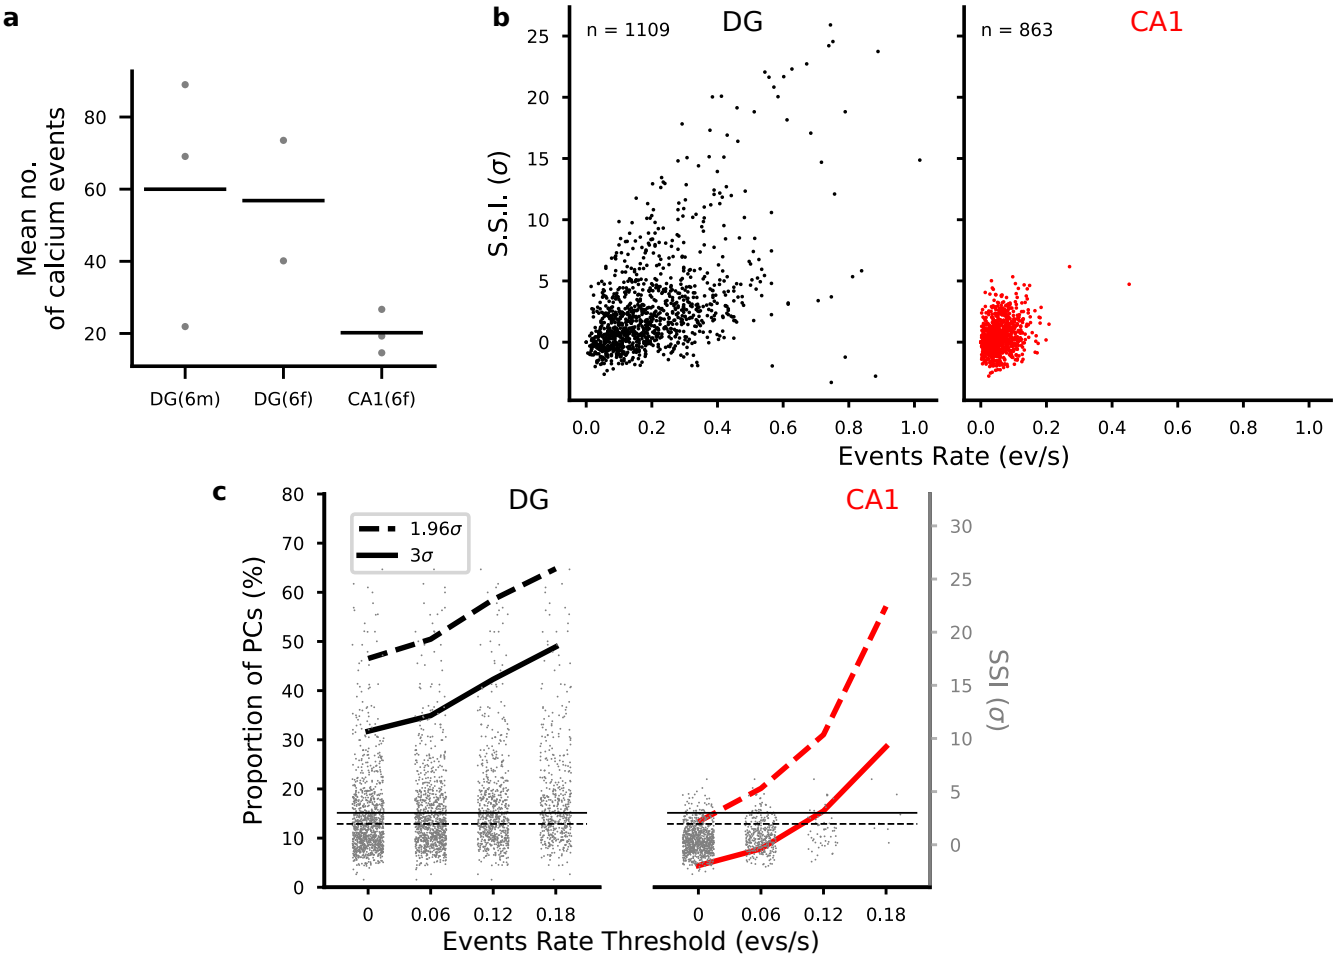

Figure S3. Reported values of overall activity levels and hence place cells ratios in CA1 and DG vary considerably across labs. This is due to a combination of factors, including different recording techniques, different statistical methods used to assess spatial tuning and differences in experimental protocol. Here we compare the significance of spatial information in the populations of CA1 and DG cells we identified and relate it to the activity levels in these areas in general, and then compare the levels of activity we observed with the existing literature. In both CA1 and DG areas, place cells were the minority of cells, however place cell ratios where very different (36% in DG, 4.2% in CA1). One reason for this difference could be in the lower activity levels we found in CA1 cells compared to DG. Place cell studies in rodents hippocampus have found higher place cell ratios when animals run through a linear track (Ziv et al., 2013). Talbot and colleagues (Talbot et al., 2018) found that 47.8% of the recorded CA1 cells were place cells but they used spatial information measure alone (Skaggs et al., 1993). The lower ratio of place cells in our report could be due to a higher sensitivity for sparsely active cells of our method with respect to standard extracellular recording approaches. Despite the likely presence of some false negatives in the single calcium event detection, even cells with a very small number of sufficiently strong calcium events are detected by the source extraction algorithm we used. Also in other recent studies, calcium imaging has revealed larger proportions of sparsely active cells suggesting these populations have been largely underestimated in electrophysiology studies (Dipoppa et al., 2018; Tang et al., 2018). In

particular, the CNMF-e algorithm uses both spatial footprints and temporal components to separate cells from background contamination and from other sources (Zhou et al., 2018). Instead, typical clustering methods for extracellular recordings may bias identification towards more densely active cells because these are more likely to be pulled out from the noise through clustering. This bias ultimately results in an overabundant representation of place cell ratios in the population. To validate our hypothesis, we first confirmed that the difference in activity levels between DG and CA1 cells was not due to the different calcium indicator used. **a**) Control that differences in activity levels between DG and CA1 are not due to the different calcium indicator used. We analyzed imaging data from granule cells in dentate gyrus after viral injection of GCaMP6f in two additional mice and found activity levels perfectly compatible with our GCaMP6m recordings, hence we conclude that differences in activity levels were not due to the calcium indicator used (**a**). Mean number of calcium events as extracted by CNMF-e across the whole session in all the DG(6m), DG(6f) and CA1(6f) mice. Each dot corresponds to one mouse. Bars correspond to mean values across mice. We then looked at the relation between activity levels in terms of event rates and spatial information (see (**b**). **b**) Scatter plot of the significance of spatial information (SSI) as a function of event rates across a session for all cell populations in DG (left) and CA1 mice (right). **c**) Proportion of place cells in DG (black, left axis) and CA1 (red) after removing cells based on their mean activity. We removed weakly active cells in both CA1 and DG populations with a varying threshold to see if we could artificially bias either population to artificially high place cell ratios. For each threshold on activity, we report the SSI for the selected cells (gray dots) and the corresponding proportion of place cells in the remaining population (black: DG; red: CA1). In our study we used a strict threshold of three standard deviations for significance (solid lines) however we also report the same proportion for a threshold of 1.96σ, corresponding to a 95% significance for spatial information as considered in some place cells studies (Kinsky et al., 2018). In line with our hypothesis, we could obtain higher place-cells ratios by choosing an appropriate threshold on activity. For example, if we didn't record any signal from cells with less than 0.12 events per second, we would have observed about 30% of place cells in CA1 and 40% of place cells in DG across all mice. The horizontal dashed and solid lines correspond to the significance thresholds for SSI for the corresponding curves (gray scale on the right-hand side of the plot). Solid lines: proportions computed using 3 standard deviations of the shuffled SSI distribution to assess whether a cell's SSI was statistically significant. Dashed lines: proportions computed on 95-th percentile threshold, i.e., using 1.96 standard deviations. The dots represent the SSIs for all cells with an activity that was larger than the given activity threshold on the x-axis. In conclusion, although differences between DG and CA1 place cell number may depend on several other factors such exposure times during habituation, our analysis suggests that the main drive for place cell counts and differences between these two areas comes from the different activity levels. In particular, a non-negligible proportion of CA1 cells are more sparsely active than previously reported and therefore CA1 place cells may have been over represented in place cells studies. We encourage researchers to adopt a more systematic assessment of the significance of the spatial information as we suggest in our work and to always report at least the proportion of place and non-place cells that are found. See also S1 for a discussion on the variability of reported rates in the literature.

Supplemental Material S4: Dependence of decoding performance from the number of cells in CA1 and DG; Related to Fig. 2

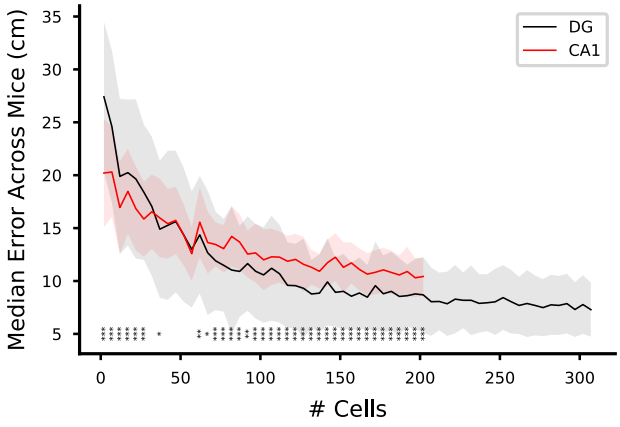

Figure S4: Dependence of decoding performance from the number of cells in CA1 and DG. For each given number of cells, a random selection of cells from the pool of identified cells is used to decode the animal's position using 10-fold cross-validation. The selection is repeated 10 times for each animal and for each number of cells. We then pooled all the results from all the animals to report the mean (solid line) and standard deviations (shaded area) of the median errors for position decoding for each given number of cells (independent samples t-test for significance, *p<0.05, **p<0.01, ***p<0.001). The maximum number of cells we used corresponds to the minimum number of cells we could record from any animal in DG or CA1.

Supplemental Material S5: Chance level decoding performance; Related to Fig. 2

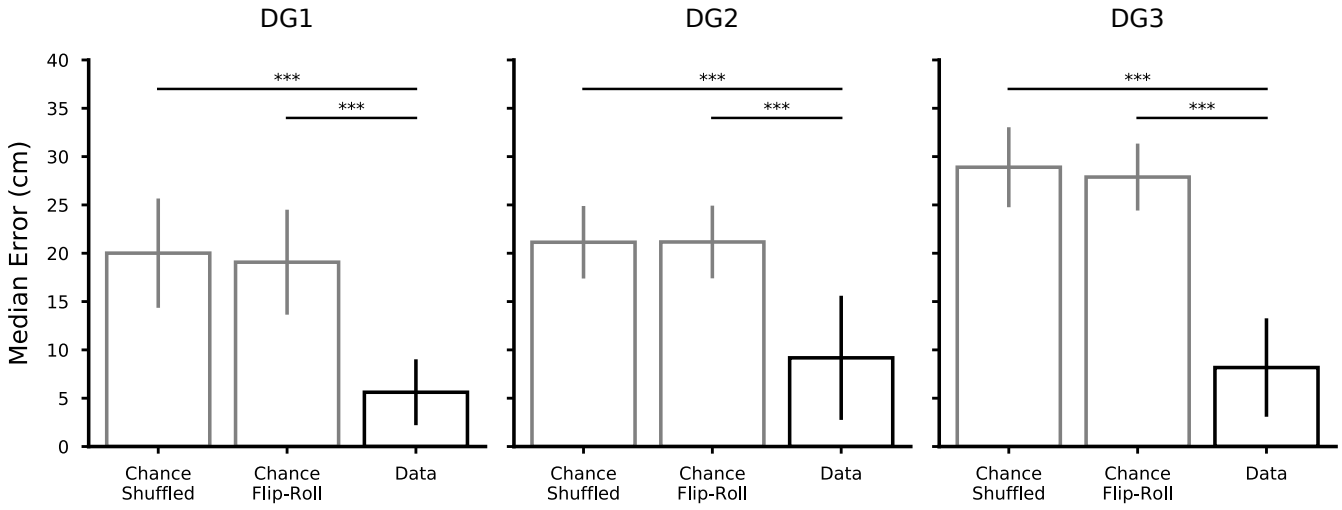

Figure S5: Chance level decoding performance. It is important to assess the performance of the decoder to a chance performance, i.e., the performance of a prediction that is not based on the recorded activities but solely on the behaviour of the mouse. The important factors that must be considered in the case of decoding continuous behaviour, are represented by the auto-correlations of the neural activities and of the behaviour. Therefore, for evaluating the chance performance one must take into account these factors in order to avoid underestimating the performance obtained by chance and therefore overestimate the ability of the decoder to extract the information under consideration. We computed the chance performance by training a position decoder on artificially manipulated data. We used two methods: the first method consisted in shuffling the calcium events in time by keeping the overall instantaneous event rate across the population constant ('chance shuffled'). The second method consisted in applying a time reversal operation to the animal x-y trajectory and then shifting it in time by a random amount. Data points falling outside the temporal window due to the shifting where reinserted from the other end of the period as in a torus.We then trained the decoder on a cross-validated manner in which the data was split into 10 chunks of contiguous data, 9 of which were used for training and the remaining one for testing the decoder. For both methods of shuffling, therefore, we obtained the 10-fold cross validated performance which we aggregated all together in each animal and applied a Mann-Whitney U non-parametric test to compare the resulting distribution to the original data (***$p<0.001$). In this figure we report the performance (mean $\pm$ s.d.) for the position decoders for all the animals for the two chance levels in grey and for the original data in black. The decoding performance is significantly above chance for both choices of chance level and the difference between the two methods is negligible. Gray bars: chance level performance (mean $\pm$ st. dev.) for shuffling calcium events in time ('chance shuffled') and for the 'flip-roll' strategy, which keeps the original correlation structure of the data intact while disrupting the association between behavior and calcium recordings. Black: decoding performance on the original data (Mann-Whitney U, ***$p<0.001$, 10-fold cross-validation)

Supplemental Material S6: Comparison of different decoding strategies; Related to Fig. 2

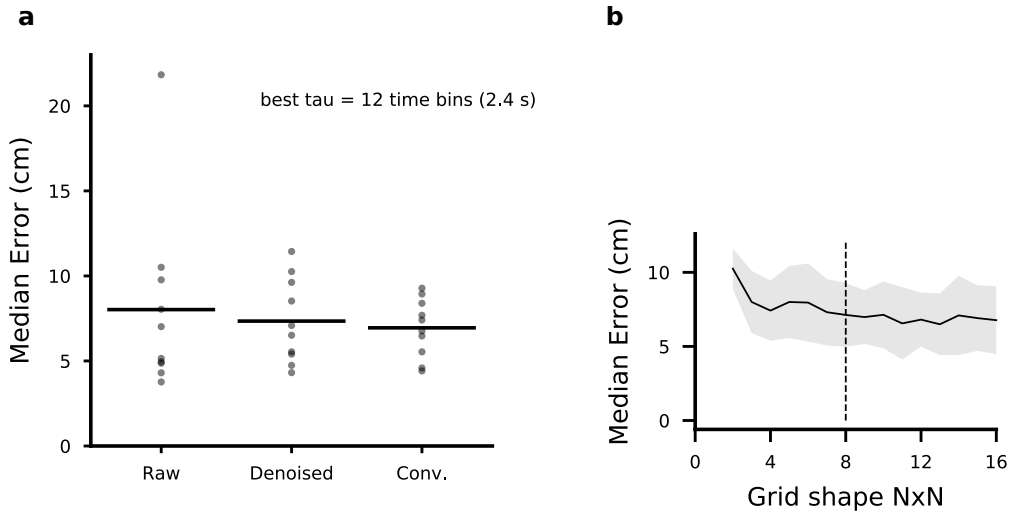

Figure S6: Comparison of different decoding strategies. To decode the animal's position and direction of motion, we used a battery of linear decoders trained on pairs of locations obtained by dividing the arena into a series of discrete locations and assigning the location label to each time bin in the data (see Methods). In this way we could use the weights of the single decoders, properly combined, to obtain an overall importance index for the cells in the population. To verify that our decoding performance did not depend on the specific choices we made, we performed the following comparisons. First, we verified that the choice of signals used for decoding did not have an impact on the position decoding. Using the convolved calcium events instead of the raw traces as described in Fig. S1 does not have a significant effect on the performance of our decoder (data for a representative mouse DG3 in **a**). We also verified that the choice of the binning to discretize the arena did not have a significant impact on the decoding performance, which allowed us to keep the same bin size on all analyzed sessions (**b**). **a**) Median error for decoding position on 10-fold cross validation for raw traces as extracted from CNMF-e, denoised traces, and convolved traces using calcium events convolved with an exponential kernel. **b**) Decoding error as a function of the number of spatial bins in which the arena is discretized (representative DG mouse, DG3, mean and st.dev. over 10-fold cross-validation). For each data point, the arena is divided into $N \times N$ squares of same size, and the size N is reported on the x-axis. The dashed line corresponds to the choice of an 8×8 grid of locations of approximately 6cm side as used in the article. Although smaller bins (larger N) might allow for a more precise determination of the position of the animal, they would also contain a smaller number of data points and therefore limit the accuracy of the decoders. The median error decreases rapidly when N is small, but then it decreases slowly and it becomes almost constant for $N > 8$.

Supplemental Material S7: Position decoding error and speed are negatively correlated; Related to Fig. 2

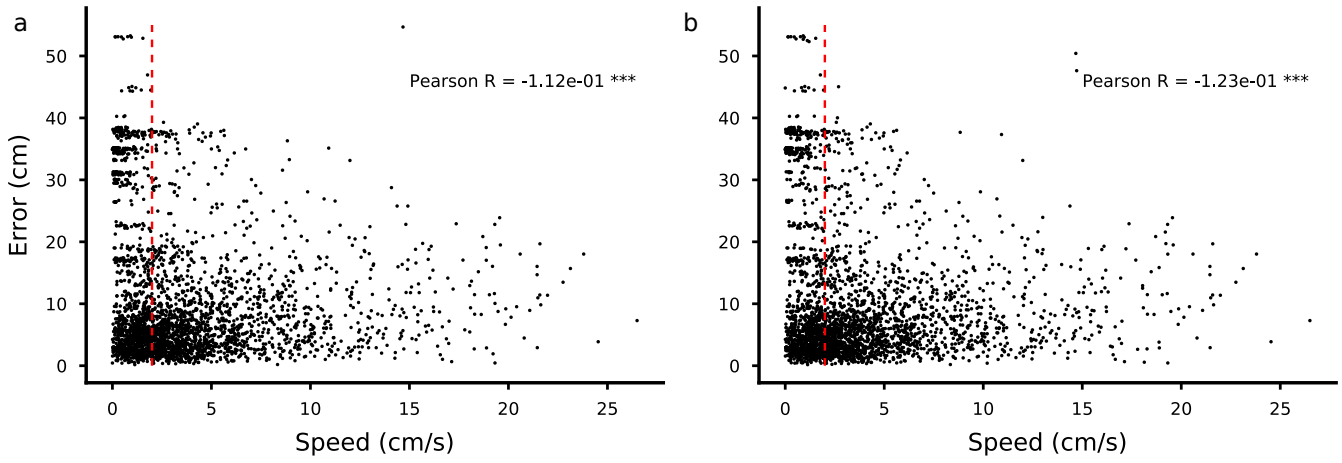

Figure S7: Position Decoding error and speed are negatively correlated. A mouse's movement and immobility states are often considered two distinct behavioral conditions, characterized by distinct neural activities, and have been therefore considered separately for analysis in the literature. In our work, we focused our analysis on the data collected during movement following a typical procedure in the hippocampal research, see for example (Leutgeb et al., 2007). We then asked whether the accuracy the decoded position was related to the speed of movement. First, we decoded the position of the animal using only the datapoints during movement for training and testing on all datapoints. **a** Decoding error at different decoding speed using only datapoints detected as movement for training. Each point corresponds to one time bin (***$p<0.001$). On the x-axis, we plot the instantaneous speed of movement in one animal (data from the dentate gyrus of a representative animal DG3). The decoded position is taken as the centre of the discrete location that is selected by the decoder in each time bin. The red vertical line corresponds to the value we used as a threshold to distinguish between movement and immobility on the training data. **b**) Same as in **a** but all the datapoints have been used for training (***$p<0.001$). To verify that the negative correlation was not due to the fact that the datapoints corresponding to immobility were not used in the training set, we repeated the procedure using all the datapoints during training. In both cases we found a significant negative Pearson-R correlation, suggesting that the animal's position was encoded more accurately during locomotion.

Supplemental Material S8: Temporal stability of the importance index; Related to Fig. 4

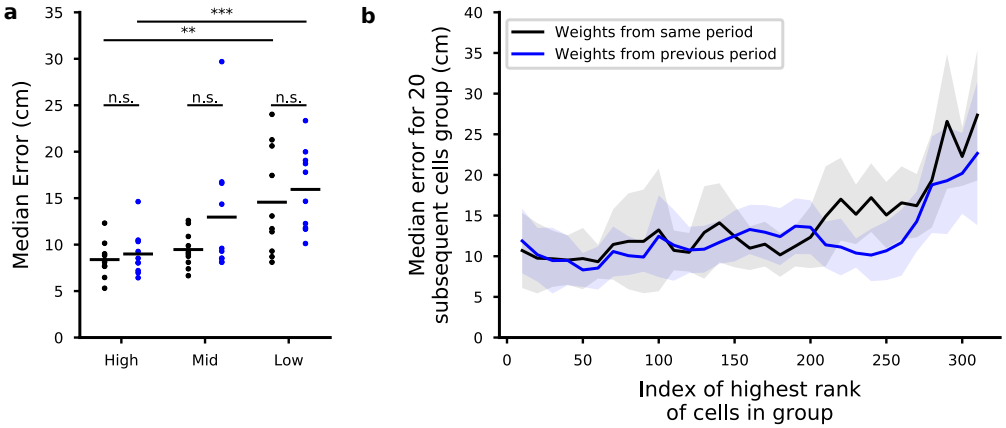

Figure S8: Temporal stability of the importance index. We introduced an importance index that quantifies the contribution of each cell to the population code. We know from the analysis in the main text that the ranking of the cells based on the importance index depends on the cells that are considered for decoding. So we know it is not an intrinsic property of the cell. However, given a population of cells that is used to estimate the importance index, it is interesting to ask whether the rank of a cell changes over time. To answer this question, we ranked the cells from the dentate gyrus using the first 10 minutes of data from the 30 minutes trials. We computed the decoding performance using the 20 best cells (ranked high), 20 cells that ranked in the middle and the 20 cells in the lowest part of the ranking. We then decoded position using each group of cells in different chunks of data of 10 minutes at the beginning of the session and in the next 10 minutes of data. If the cells' ranking changed, we'd expect a different performance for each group of cells since they were defined using the ranking in the first 10 minutes. Instead, we saw no significant changes in decoding performance for the high, middle and low-ranked cells. This indicates that the importance of each cell is relatively stable over time, at least when the decoding performance is considered. **a**) Decoding performance for three groups of cells 20 cells divided by their ranking in importance index. Black: decoding performance for 10-fold cross-validation within the first 10 minutes of data. Blue: decoding performance in the following 10 minutes of data after training on the first 10 minutes (Mann-Whitney U, n.s. $p>0.05$, **$p<0.01$, ***$p<0.001$). **b**) Decoding performance for subsequent groups of 20 cells ranked by their importance index. Black: cells were ranked using the importance index computed on the first 10 minutes of data. Blue: cells were ranked based on the weights computed on the next 10 minutes of data. We applied a procedure like that of Fig.4 of the main text to better assess whether the ranking in one time-period could be used to rank cells in another time period. We first computed the decoding performance in the first 10 minutes of the trial using subsequent groups of cells ranked by their importance index computed in the same time window. We then compared these results with the ones obtained by decoding position during the next 10 minutes of the trial. The two plots largely overlaps, therefore the importance indices obtained in one temporal window can be used to select the important cells in another temporal window.

Supplemental Material S9: Importance index and spatial information; Related to Fig. 5

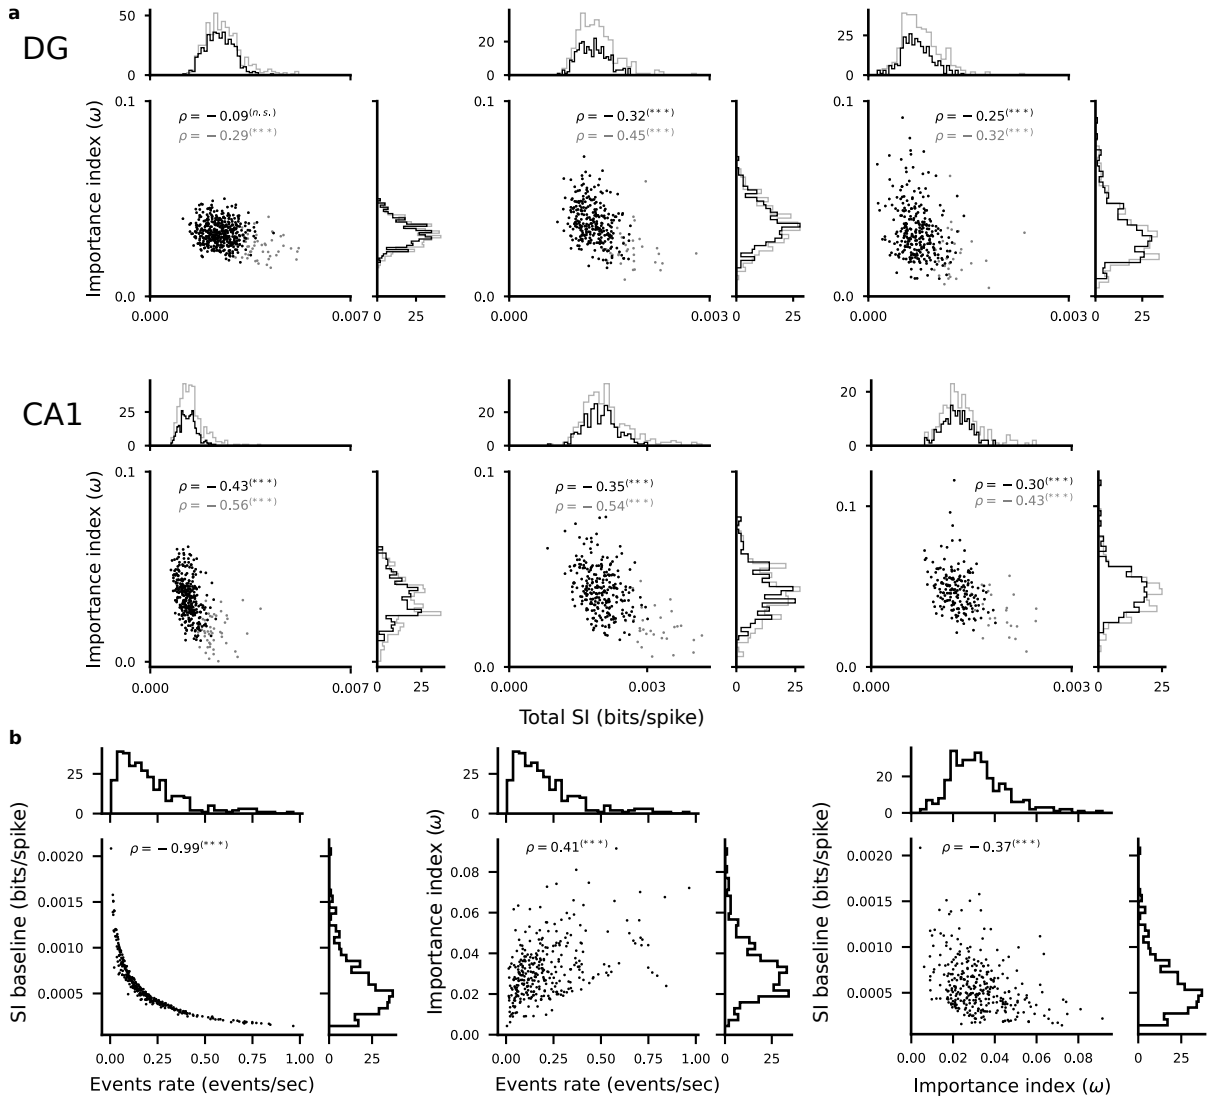

Figure S9: Importance index and spatial information. Limited data, such as low firing rate, introduce a systematic bias in the estimation of the information content in neural activities (Panzeri et al., 2007). Our results show that a more sensible choice is the one of the SSI, a measure of the strength of the spatial information after compensating for this bias. Furthermore the importance index was weakly correlated the SSI. We also showed that importance index and SSI are correlated. Here we additionally show that importance index and spatial information have a low and even negative correlation factor (Pearson-R correlation, ***p<0.001). **a)** Relation between importance index and spatial information for all recorded animals (grey: all cells, black: cells with more than 10 calcium events identified, Pearson-R reported in each plot, ***p<0.001). Each dot in the plots represents one cell. The histograms for each of the quantities are shown on the sides of each plot. This result is further evidence that the spatial information alone may be a misleading factor in estimating the contribution of a cell to encoding position if not validated through an assessment of its statistical significance. **b)** Left: Relation between the baseline spatial information per cell (obtained by shuffling events in time for each cells) and events rates (left), importance index and events rate (middle) and spatial information baseline and importance index (right). In each panel,

each dot corresponds to one cell in a representative animal (DG3). The spatial information baseline is computed as the mean spatial information over the shuffled distribution. Spearman-ρ correlation values and significance are reported in each plot. The baseline and the importance index are negatively correlated due to their relation to event rate. Intuitively, the negative correlation between spatial information baseline and importance index is due to the fact that low activity introduces spurious stimulus-dependent differences in the response probabilities used to compute information (Panzeri et al., 2007; Kelemen and Fenton, 2010). These make the neuron seem more informative (high spatial information) despite its low importance in the population for position encoding.

Supplemental Material S10: Decoding performance using only place cells; Related to Fig. 4

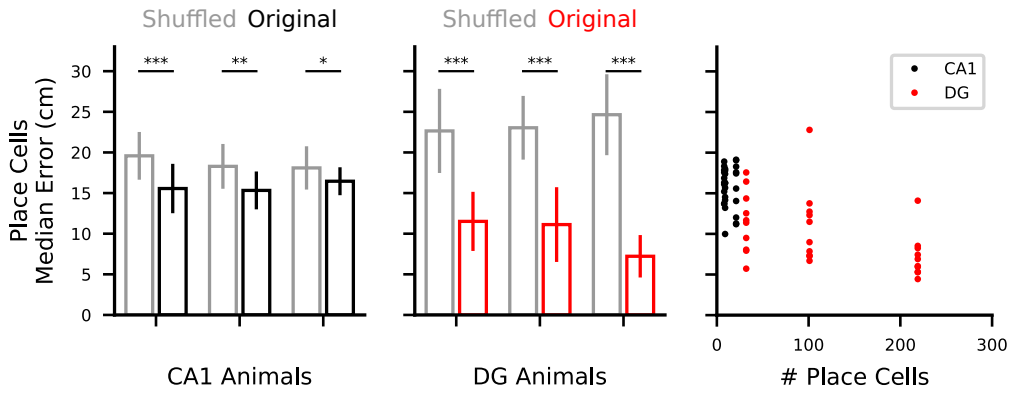

Figure S10: Decoding performance using only place cells. Using only place cells, the decoding performance is close to the best performance we obtained in DG (left, chance performance in grey) but rather limited in CA1 (middle panel). This is because the decoding performance depends on the number of place cells used, which is limited in CA1 (right). (Mann-Whitney U, *p<0.01, **p<0.005, ***p<0.001.) Left and Middle: Median decoding position error using only place cells for CA1 (left) and DG (middle) animals (each bar represents mean and st.dev.; Mann-Whitney U test, *p<0.05, **p<0.005, ***p<0.001, see Methods). Right: Median decoding error for all CA1 (black) and DG (red) animals for each 10-fold cross-validation within the 10 minutes sessions aligned to the corresponding number of cells isolated in the corresponding FOV.

Supplemental Material S11: Control for correlation between movement direction and position; Related to Fig. 7

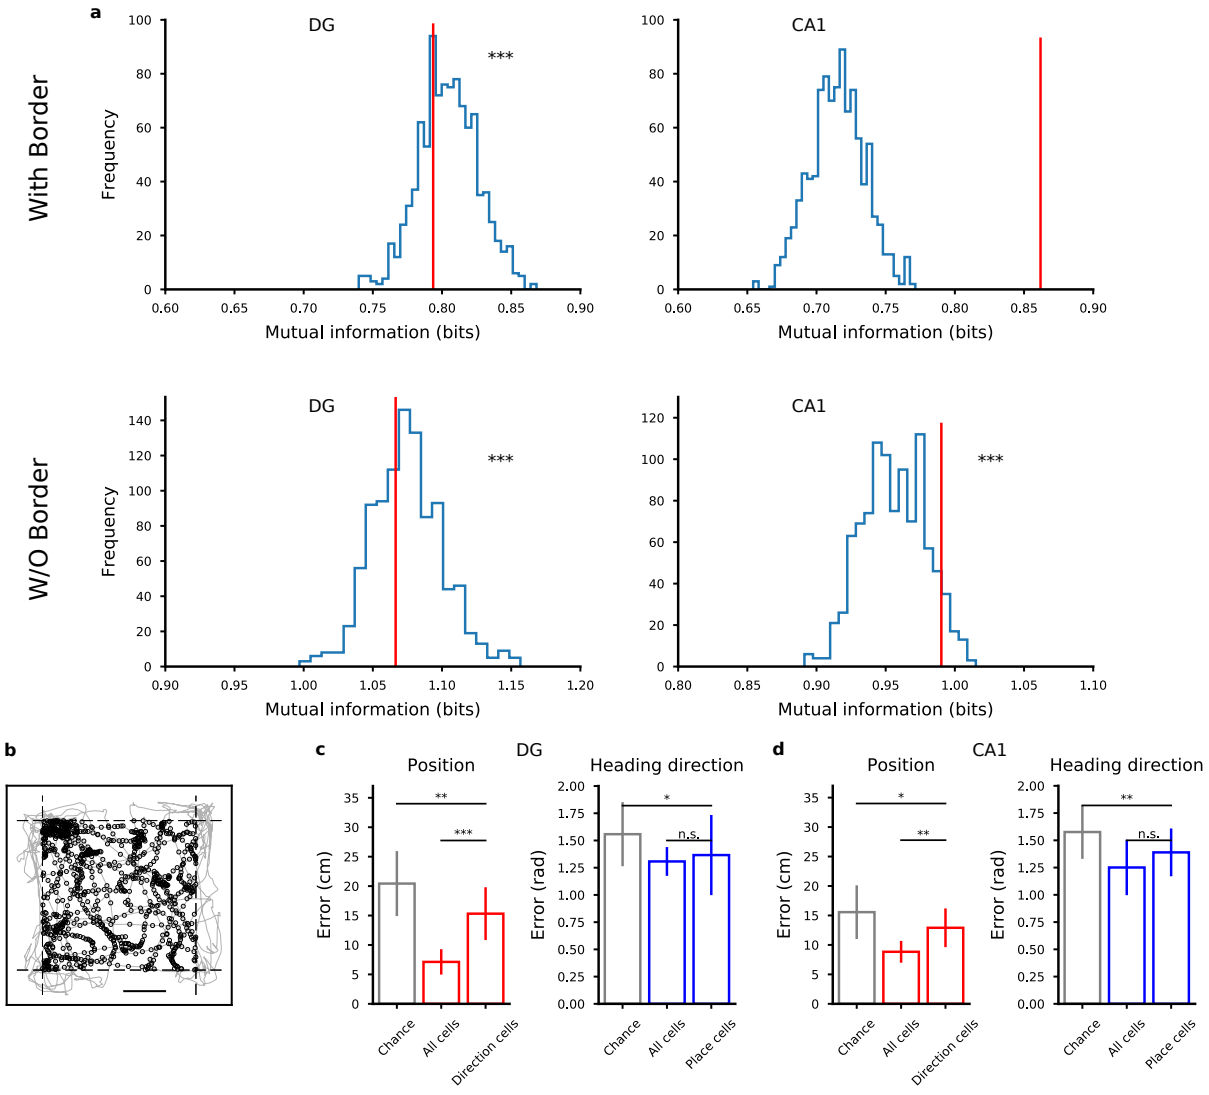

Figure S11: Control for correlation between movement direction and position. We showed in the main text that we can separately decode position and direction of motion and that the neural code for these two variables appears distributed across the neural population (Fig. 7). Our strategy was to show that the important cells for decoding position could be used to decode the direction of motion. However, it could be that the reason why direction could be decoded using the important cells for position was that direction and position were correlated. Indeed, when the animal is very close to one the border of the arena, the direction of movement can only assume values that point towards the inside of it. The decoder could in principle use this information to decode direction of motion from the important cells for position, therefore confounding our results. To exclude this possibility, we computed the mutual information between direction of movement and position and compared it to the distribution obtained with 500 shufflings of data from the same session from the animal with the most homogeneous coverage of the environment. **a**) Mutual information between position and direction of motion for shuffled data in the same session (blue histogram) and real data (red vertical line). Top panels: original data with positions close to the arena walls. Bottom: filtered data as in **b**. DG in top left panel, CA1 in top right panel, same animals

as in Fig. 7. In the shuffled data, position was time reversed and shifted similarly to what was done for evaluating the chance performance (see also Methods in the main text). The mutual information in the real data (red line) lies within the distribution of mutual information for shuffled data (blue histogram, one sample t-test, ***$p<0.001$) for DG, hence the knowledge of one variable is not enough to predict the value of the other variable, but not for CA1. We attribute this difference to the effect of the borders in the CA1 experiments, in which a rectangular arena of half the size than the one for DG mice was used. **b**) To further verify that the correlation between direction of motion and position did not explain the results of Fig. 7b, we repeated the decoding analysis after removing from the data all the positions that were recorded in the spatial bins closest to the walls of the arena. Here we show the trajectory of a representative DG mouse (DG3). Positions close to the walls of the arena, i.e., beyond the dashed lines, were eliminated from the data. Grey: mouse trajectory. Black dots: positions included in the analysis. Same DG mouse as in Fig. 7. After this manipulation, the mutual information between position and head direction becomes indistinguishable from the chance distribution in both DG and CA1 (bottom panels in **a**) and we were still able to decode position from the most important cells for direction and vice versa, as we did for the original data in the text, for both DG and CA1 mice (see **c** and **d**). **c**) Decoding error for important cells for position (left) and direction (right) from a representative DG mouse, same procedure as in Fig. 7b in the main text but after removing border data as in b. **d**) Same as in **c** but for the CA1 mouse of Fig. 7b.

Supplemental Material S12: Temporal stability of the decoder; Related to Fig. 2

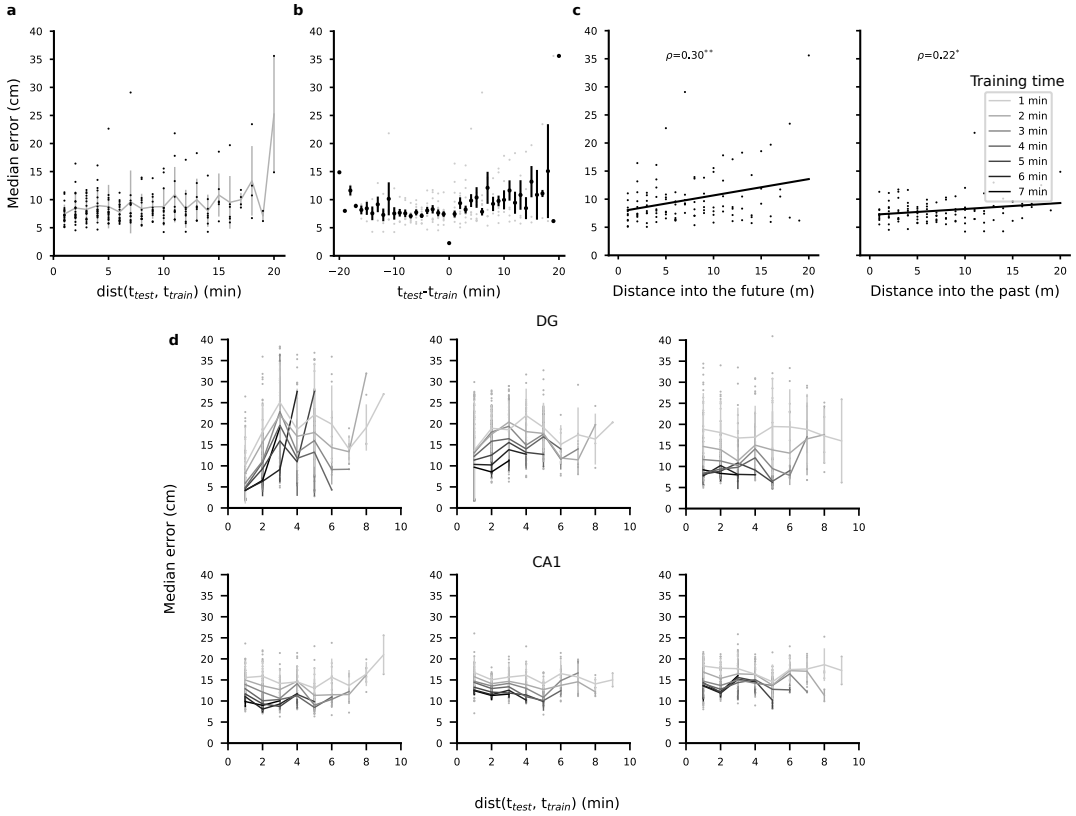

Figure S12: Temporal stability of the decoder. We asked whether the representations for position changed within a recording session. To do this, we used the 30 minutes trial available in our data (animal DG3, data from the dentate gyrus). We trained a position decoder on a shifting window of 10 minutes of data and tested in on a shifting window of 1 minute of data. **a**) Position decoding performance as a function of the distance in time between training set and test set. Each dot corresponds to one choice of training and test set, the grey line joins the mean values on pooled data for each bin of distance in minutes, bars are standard deviations within the binned distances. The performance of the decoder is remarkably strong and stable for a period of time extending to the available 20 minutes in the trial, although with some slow degradation towards the more distant time windows (**a-c**). This result shows that position is encoded with strong accuracy on the populations of cells whose coding properties extend long in the trial and it seems compatible with results in literature showing that the representation of position is stable across sessions despite the large degree of variability on the tuning properties of single cells (Ziv et al., 2013). **b**) Data as in (**a**) but datapoints for which the test precedes the training set are on the negative side of the x-axis. We looked for evidence of a difference between testing the decoder in past time periods and future time periods with respect to the training period. To determine whether decoding data temporally preceding the training data (past) was different from decoding following data (future), we linearly interpolated the decoding performance and obtained two weakly significant Pearson correlation factors of 0.3 for the future (**p<0.01) and 0.22 for the past (*p<0.05). Given the weakness of our statistical test we couldn't detect a difference between these two trends. Therefore, the performance on the test data does not depend on whether the decoder was trained on a past or on a future interval. **c**) Data as in (**b**) where the points have been fitted with a linear regression for test data that is in the future with respect to training data (left) or in the past (right). Reported correlation values correspond to the Pearson's R correlation (*p<0.05,

p<0.01). **d) We compared the stability of the decoder in DG and CA1 animals by repeating the analysis within 10 minutes of the session for all animals. For short training times (down to 1 minute) the decoding error was relatively stable on most animals but was much lower that the best decoding performance we obtained for those animals. We then increased the training time to verify that this was not due to a low temporal stability of the code but rather to the small size of the training set. In all animals the performance drastically improved with training set size. These results further suggest that both in DG and in CA1 the code for space is relatively stable within a session and further studies will further investigate this important aspect of the spatial code. Top: DG animals. Bottom: CA1 animals. Each panel corresponds to one animal. Each line in different shades of gray corresponds to a different choice of training time (test time is always 1 minute). Each dot represents the median error for one choice of training and test time. The decoding performance for different test periods (y-axis) are sorted according to the temporal distance between the test period and the closest datapoint in the training period (x-axis), regardless of whether the test period preceded or followed the training period (t_{test} is the time of the first time bin of the test period, analogously for t_{train}). In each one-minute bin, mean and st.dev. are computed and plotted.

Supplemental Material S13: The information about position is distributed: distributions of importance indices in DG and in CA1; Related to Fig. 5

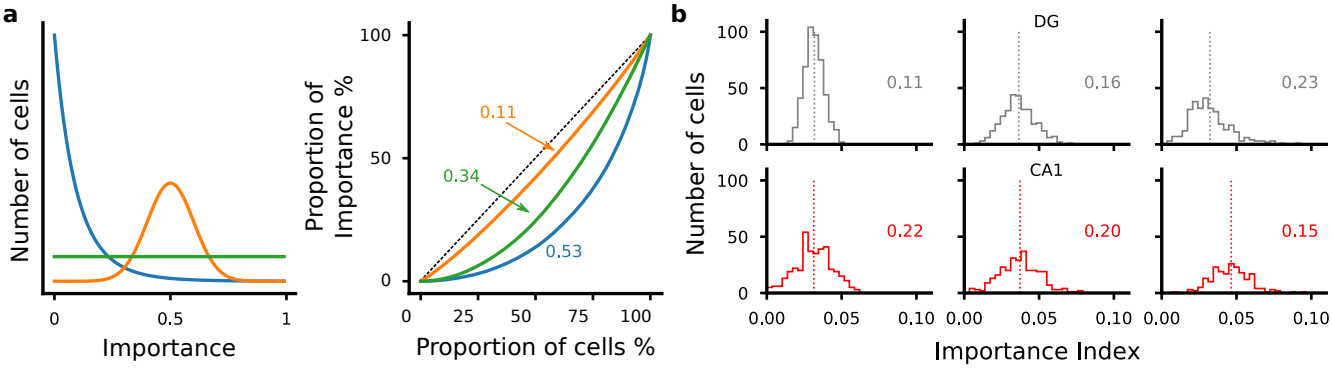

Figure S13: The information about position is distributed. One way to compare how populations of cells represent information is to compare the distribution of their importance indices. We computed the Gini index on the distribution of importance indices to assess whether the information about position was distributed across multiple cells or instead carried by a small minority. The Gini coefficient (Gini, 1912) is a measure of statistical dispersion among values of a certain frequency distribution. It is often used as a measure of inequality of income levels. Low values correspond to a large degree of equality, i.e., people share large proportions of the total wealth of the population. High values often identify situations in which wealth is concentrated on a small number of people. In the same spirit, we used the Gini coefficient to evaluate how the importance index was distributed across cells in our recordings. **a**) To illustrate the dependence of the Gini value from the underlying distribution, we first describe how different distributions result in different Gini indices. We considered three different distributions, from high inequality to low inequality. Here we show the distributions (left) and corresponding Lorenz curves (right) for simulated data: a uniform distribution (green), a distribution with low inequality (orange), as in the case of a distributed code, and one with strong inequality (blue) as in the case of a specialized code with few important cells and many unimportant cells. The first one is the Pareto distribution, a power-low distribution commonly used to describe social phenomena, including distribution of wealth. In this distribution, importance is concentrated on a small fraction of cells whereas the vast majority of cells have a low importance, as in a specialized code (left panel, in blue). We also considered a uniform distribution of importance whereby a cell importance is determined by chance (left panel, in green). Lastly, we considered a Gaussian distribution centered around a middle value. This case results corresponds to a highly distributed code since the vast majority of cells shares a similar amount of importance (left panel, in orange). The first step for computing the Gini index consists in evaluating the Lorenz curve, which represents what proportion of cells shares a certain proportion of importance. The area under the Lorenz curve is then compared to an ideal situation whereby all cells have the exact same amount of importance, resulting in a straight Lorenz curve, and hence in an area under the curve of 1/2 (dashed black line). Finally the Gini index is computed as the proportion of the area defined by this ideal distribution that is not covered by the Lorenz curve of the data. In the right panel we plot the Lorenz curves for the three distributions we considered and their respective Gini index values. As expected from intuition, the most equal distribution is the Gaussian distribution (right). This distribution results in an extremely low value of Gini (colored inset values on the right). **b**) The Gini index and the distribution of the importance indexes across neurons. We computed the Gini index on the values of importance indices for both DG and CA1 and found low values of Gini (between 0.11 and 0.23), indicating that the code is highly distributed.

Histograms of importance indices from the data across all DG (top) and CA1 animals (bottom). Gini index numbers are reported in each plot.

Supplemental Material S14: Procedure to destroy correlations and controls; Related to Fig. 8

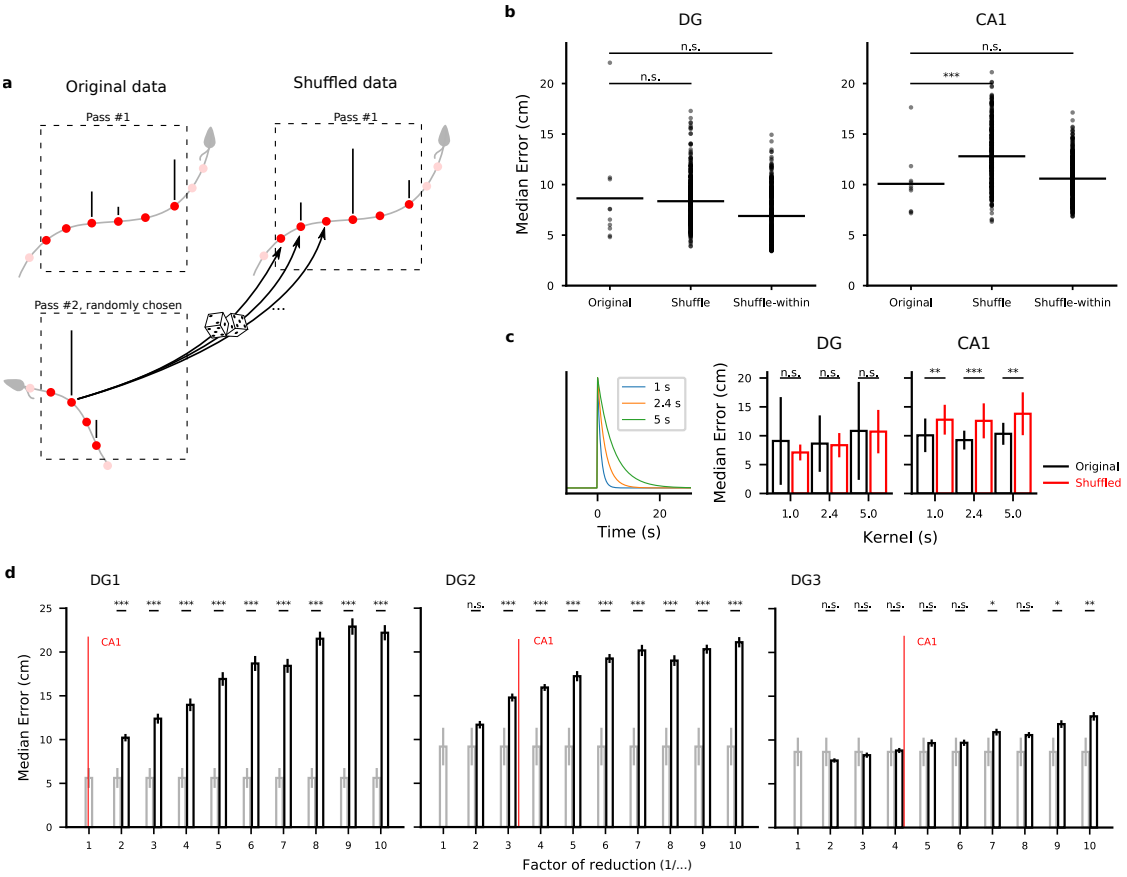

Figure S14: Procedure to destroy correlations and controls. **a**) Schematic description of the procedure for destroying correlations in one discrete location in the arena. The red dots correspond to the time bins of the recorded calcium events along the mouse trajectory (gray line) through the area of interest (dashed line). Time bins in which an event was recorded are identified with a tick mark on top of the red mark. Light red marks correspond to datapoints recorded before the animal enters and after it exits the area and are not considered. For each location, we focused on all the passes of the animal through that location and identified the entrance and exit times of each one. Then for each pass, we substituted the recorded calcium events in that pass with the ones of an other randomly chosen pass. To account for the different number of time bins in the different passes, we randomly sampled from activity from the second pass. The procedure showed an effect in CA1 data but not in DG (see **b**). This suggests different coding properties as discussed in the text. **b**) Decoding error for original data, shuffled data as in a and shuffled data when the new data was generated from the same pass to control for artifacts due to sampling. Left, DG. Right, CA1 (Mann-Whitney U test, ***$p<0.001$). To verify that random sampling did not introduce artifacts that would affect decoding performance, we verified that when we sampled from data within the same pass the decoding performance was not significantly affected ('shuffle-within'). After this procedure, we were confident that any effect we would observe was only due to having destroyed correlations among neurons. When we shuffled data within the same pass, we saw no effects in either CA1 nor DG data. **c**) The effects of destroying correlations do not depend on the choice of the kernel time scale that was used to convolve the sequences of events. Left: shape of the kernel function used to convolve the calcium

event data. In the main text we used 2.4 s. Right: decoding error for the original and the shuffled data in representative DG (left) and CA1 (right) mice for different choices of kernels, as in Fig.8 in the main text (Mann-Whitney U test, **p<0.005, ***p<0.001). **d**) The effects of destroying correlations depend on sparsity. Sparsity was induced by randomly selecting a given proportion of detected calcium events. We report the performance on shuffle data (black bars) compared to the original decoding performance (grey bars) for all the DG animals at different levels of sparsity. The red vertical line corresponds to the equivalent factor to obtain similar activity levels to those of the CA1 animals, for reference (Mann-Whitney U test, *p<0.05, **p<0.005, ***p<0.001). When we artificially reduced the number of events in our DG mice we observed a significant decrease of performance for increasing sparsity, however in two of the three analyzed animals the reduction was not significant at a level of sparsity that was comparable to CA1.

Supplemental Material S15: The effects of dimensionality on neural correlations and decoding performance; Related to Fig. 8

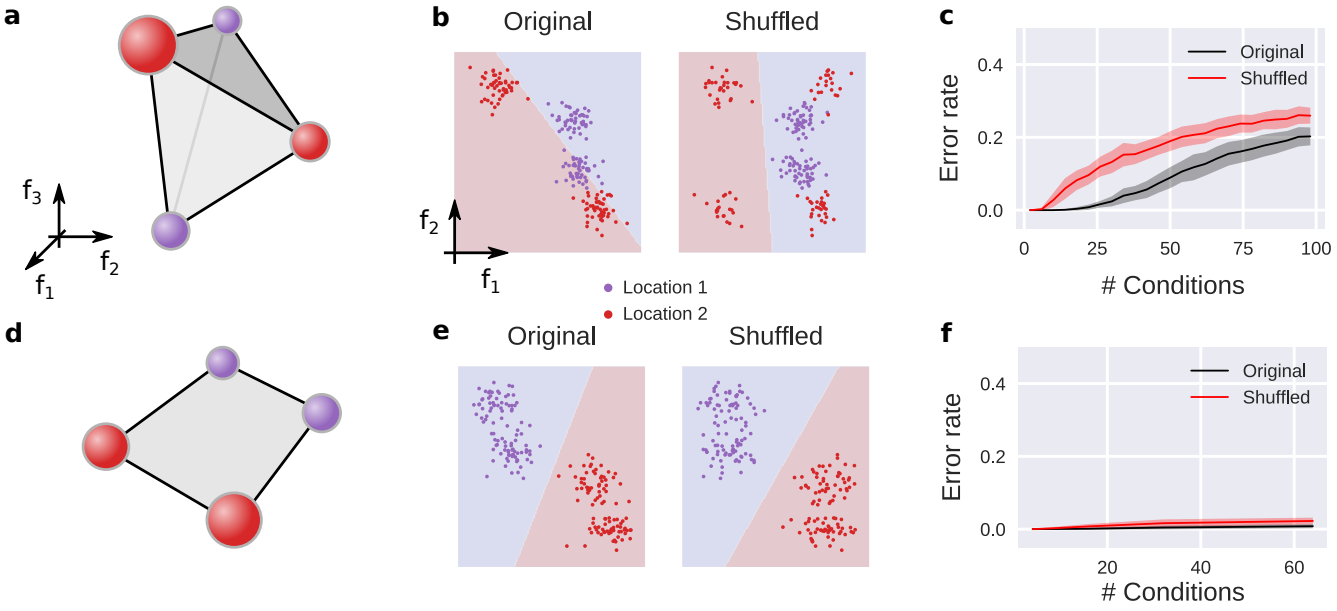

Figure S15: The effects of dimensionality on neural correlations and decoding performance. In our work we suggested that, at least in CA1, the neural correlations should not be considered as noise but rather as a reliable signal. We showed that destroying the neural correlations had an effect on the ability of our decoder to predict the animal position from CA1 activities. The effect was smaller in DG. In general, the effect of destroying correlations on decoding performance depends on the geometry of the neural representations (Abbott and Dayan, 1999), i.e., how the patterns corresponding to the different combinations of behavioral variables are distributed in the space of population activities. How do the effects of destroying correlations depend on the geometry of the neural representations? Here we studied in simulations a few cases to understand under what assumptions we should expect a disruption of the performance when the correlations are destroyed. We started by considering neural representations in the space of population activities. These are determined by different combinations of discrete behavioral variables, each of which corresponds to an experimental condition. In the experiments, some of the behavioral variables would be under direct control, such as the discretized animal's position or movement direction, while others are not. For each condition we generated a prototype pattern of population activity and a cloud of points around it by adding isotropic Gaussian noise of zero mean and unit variance (see S15). This cloud of points thus represents a set of recordings for that condition, similar to the ones generated by multiple passes of the animal through a given location in the environment as in Fig. 3 in the main text. Following our hypothesis, we generated many such clouds in the space of neural activities for the various combinations of hypothetical behavioral variables the neural activity is subject to and assigned half of them to one class, corresponding for example to one location, and half of them to another class, corresponding to a second location. The fact that the representations of each location consists of multiple conditions, naturally induces correlations among neurons in the population. This is because the activity of one neuron is tightly coupled to the activity of other neurons as prescribed by the patterns corresponding to that location. To study the effects of destroying correlations between, we first considered the cross-validated ability of a linear decoder to correctly classify the two hypothetical locations. Then, we shuffled the data in a way similar to what we described in the Methods to destroy the correlations among

neurons while keeping the information on location. Briefly, for each location, we chose one observed level of activation in that location for each neuron independently. This effectively violates the coupling imposed by the patterns of the conditions representing one location and therefore destroys correlations among neurons. Importantly, while this manipulation effectively destroys the correlations among neurons, it does not affect the statistics of the activity for each neuron with respect to each location, i.e., the spatial tuning of the neurons is not affected. We repeated this procedure multiple times to generate a new dataset of data without neural correlations. After this manipulation we trained again our decoder and compared the decoding performance with the one computed on original data, as we did for the neural recordings in the main text. To being our study of how the geometry of the neural correlations affects decoding, we considered two common scenarios observed in neural data (Rigotti et al., 2013; Bernardi et al., 2018), one in which the neural representations are 'unstructured' and one in which they are 'structured'. The first situation can be generated by randomly distributing the different conditions in the space of neural activities, as in **a**. For the second situation we considered abstract variables, i.e., the conditions were positioned at the corners of a low-dimensional hypercube randomly rotated in the space of neural activities as in **d** (Bernardi et al., 2018). To visually describe the situation, in **b** and **e** we show these two scenarios in the simple case of two neurons. We further restrict our description to the case of two conditions per location, which could correspond for instance to two directions of motion in the data. The colored regions separated by a straight line represent the decision function of a linear classifier trained on cross-validated data, hence the red points lying on the red regions are correctly classified as are the purple points lying on the blue region. Vice versa, the points lying on a region of different color are incorrectly classified. In the case of unstructured data, the four clouds of points appear randomly distributed (**b**) and, in this example, they are linearly separable. The multiple conditions impose correlations between the two neurons, in a way similar to what we discussed in Fig. 3 in the main text. Indeed, after destroying the correlations, the two location are no longer linearly separable and therefore the performance of the decoder can only worsen. The situation in the high-dimensional space of the population activities doesn't have such a straight forward graphical interpretation and so it is not straight forward to predict what would happen in the high-dimensional case. Thus, we performed simulations in the high-dimensional case (50 neurons) and compared the decoder performance before and after destroying the correlations for increasing number of conditions (**c**). We found that the effect of the correlations increases with the number of conditions (**c**), it reaches a maximum and then it decreases again for higher numbers of conditions. Moreover, when the number of conditions becomes comparable to the maximum number of neurons (50 in our simulations), the two locations may be non-linearly separable also before destroying correlations, depending on the specific random arrangement, and so the error rate increases also for the original data. In the case of structured neural representations, the situation is different (**d**, **e**, **f**). For this scenario, we generated the patterns in a way that follows abstraction as defined in (Bernardi et al., 2018) (**d**). As in the previous case, we first looked at a low-dimensional case of two neurons and two conditions per location (**e**). In this case the effect of the correlations is much weaker, as it is shown by the fact that also after destroying correlations the representations of the two locations are still linearly separable. To numerically explore the high-dimensional case, we performed similar simulations as before and varied the number of conditions while preserving the structure of abstraction (**e**). We observed that even for a large number of encoded variables the effect of destroying correlations was negligible in this case. Taken together, these results show that the geometry of the neural representations imposes neural correlations can be beneficial for decoding performance. In some of the cases we examined we also found that the effect can be negligible. It is important to note that there can be other scenarios that we did not consider here but even with the few simple cases we described we found several conditions that seem compatible with our experimental data. Furthermore, it is interesting to note that the type of neural correlations that we considered in these

models are not induced by a non-negligible covariance in the direction of the noise of neural activities. Rather they are induced by the presence of relatively small regions of the space of population activities that correspond to reliable representations to experimental conditions. It will the subject of further studies to investigate what classes of models can lead to such sparse representations. The impact of destroying neural correlations depends on the dimensionality of the neural representations and affects high dimensional representations more than low-dimensional ones. **b**) Two-dimensional projection of the activities recorded ad different times in each of the four conditions in **a**. The line separating the red and purple regions of the space corresponds to a linear separation of the red and purple locations. Left: original data. Right: datapoints obtained after destroying correlations as explained in the Methods. **c**) Decoding position, i.e., decoding performance for discriminating red from purple conditions, as a function of the number of conditions for original and shuffled data. Half of the conditions are randomly assigned to each class. **d-f**) Same as in **a-c** but the conditions are arranged on the vertices of an hypercube as in **d**.

Supplemental Material S16: The effects of shuffling on pairwise correlations; Related to Fig. 8

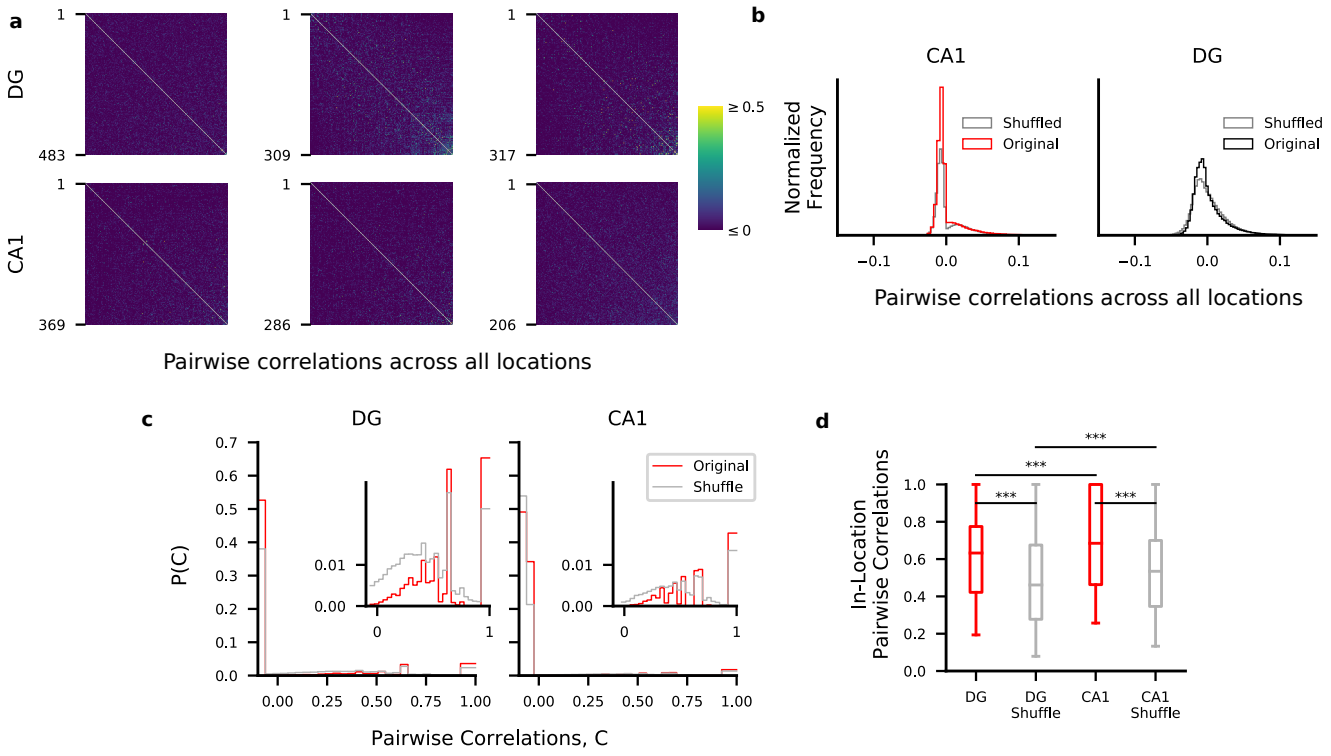

Figure S16: The effects of shuffling on pairwise correlations. Our procedure for destroying correlations shows a significant effect in CA1. To further characterize these effects, we studied the pairwise correlations before and after the shuffling. Although pairwise correlations only partially reflect the structure of correlations that is disrupted by our procedure for destroying correlations, because they include those induced by the similarity of tuning among cells, one may wonder if they would be sufficient to describe the situation. However, it is important to notice that pairwise correlations include those imposed by the tuning profiles of the cells (often referred to as signal correlations). In contrast, our procedure was specifically introduced to study the effect of correlations that appear on single passes through one location (often called noise correlations). When we destroy correlations, therefore, we disrupt those induced by the influence of variables other than position, i.e., a combination of noise and other potential signals that influence the cell activity such as movement direction. We show in Fig. 8 of the main text that our procedure does not disrupt spatial tuning of the cells. It could be therefore misleading to draw conclusions from the pairwise correlations alone about how they influence the conjunctive representations of other variables. Rather, we would like to make the point that by looking at the effects of destroying correlations while keeping the spatial tuning of the cells we can select hypothesis on the relations between the encoded variables, as we show in Fig. S14. Our analysis suggests that the structure of the pairwise correlations is only partially informative and that is important to use more sophisticated quantities to investigate the role of neural correlations. **a**) Here we show the pairwise correlation matrices for DG (top) and CA1 (bottom) animals computed using the calcium event time series across the entire session (three FOVs for each area). The values of the correlations are overall low in both regions, with medians around 0.1-0.2 but long positive tails (see the distribution in **b**, absolute values). **b**) Histograms of pair-wise correlations across all cells, computed across the entire session, pooled across mice (3 in DG and 3 in CA1) before

(red and black) and after applying the shuffling procedure to destroy the correlations (gray histograms; see Methods). We did notice a difference in the overall distribution of pairwise correlation values in the two different areas. In particular, CA1 distributions seem to show a stronger peak for negative values and tend to have a smaller frequency of values around 0. The distributions didn't qualitatively change after destroying the correlations through shuffling in either region. **c** To investigate correlations with finer detail, we studied their marginal distributions within single locations in the arena (histograms for all DG and CA1 cells combined across animals). Insets show a magnification of the y-axis. The peak of the distribution of the correlations across pairs of neurons was negative and small in magnitude since in the vast majority of passes only a few cells are active at any point in time. However, there are several pairs of neurons with a relatively strong correlation, as shown in the insets. As expected, the peak of the distribution was concentrated within a short interval of small negative values since in the vast majority of passes only a few cells would be active at any point in time. We then focused our attention at the positive side of the distribution (inset), since it is the largest contribution to correlations (in absolute terms). In line with our main results, we found higher positive correlations in CA1 than in DG (Mann-Whitney U, ***$p<0.001$). **d**) In both regions pairwise correlations are affected by our procedure. Here we show the distribution of the positive noise (in-location) correlations. In line with our main results, we found higher positive correlations in CA1 than in DG. Whisker plots show median, quart-percentile boxes, 5th and 95th percentile whiskers and Mann-Whitney U statistical tests for significantly different distributions (***$p<0.001$). Altogether, these results further confirm that although some correlations exist within DG, pairwise correlations only partially explain the effects on decoding that are caused by our shuffling procedure.

Supplemental Material S17: Importance and decoding performance for position; Related to Fig. 4

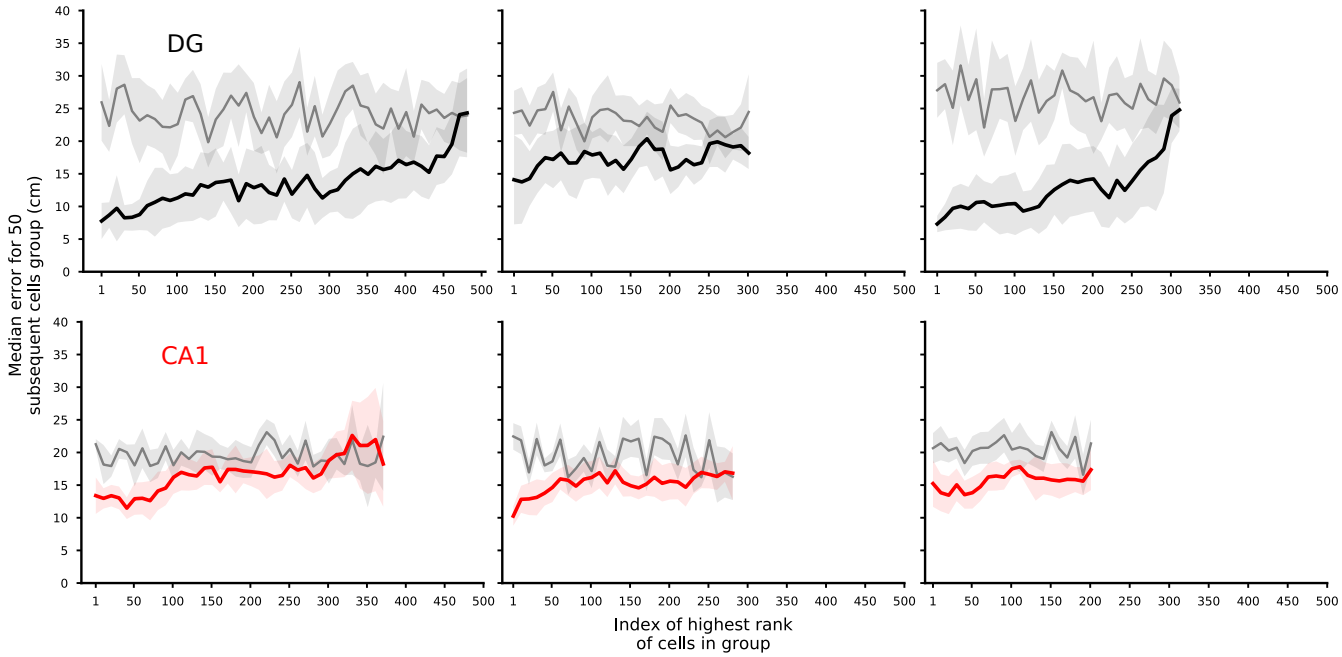

Figure S17: Importance and decoding performance for position. Decoding performance for cell ranking. Top: DG animals. Bottom: CA1. One animal per panel. See Fig. 4a in main text.

Supplemental Material S18: Importance and decoding performance for direction; Related to Fig. 6

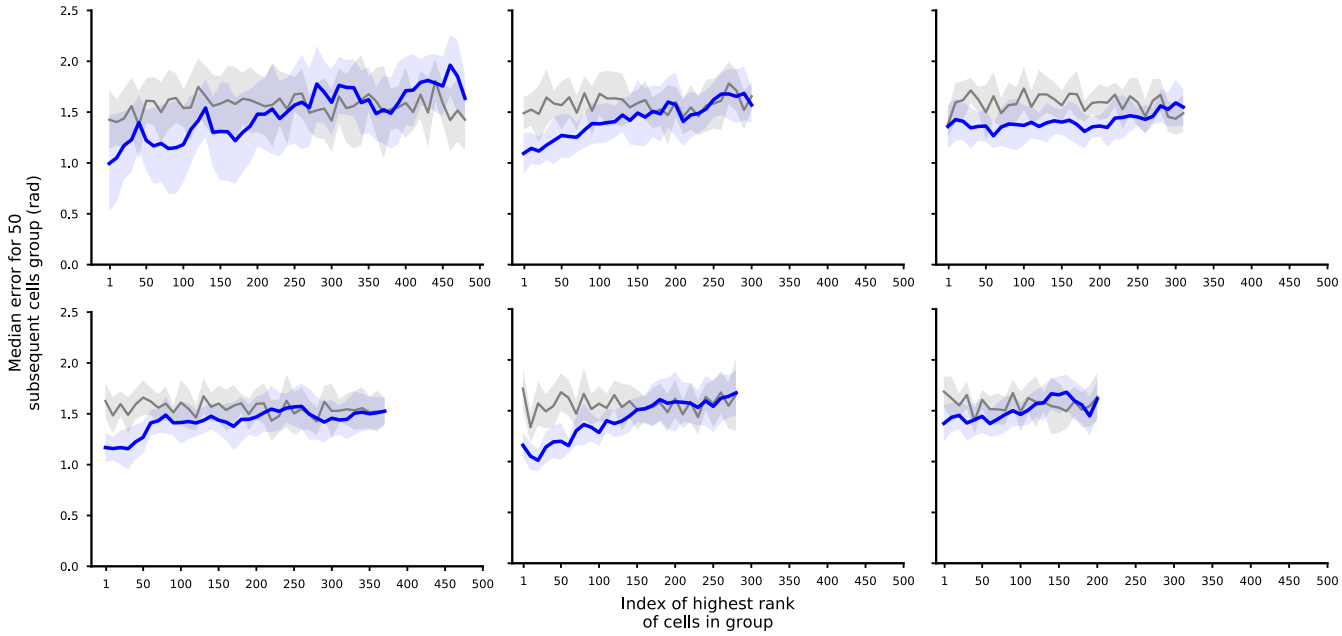

Figure S18: Importance and decoding performance for direction. Decoding direction performance for cell ranking. Top: DG animals. Bottom: CA1. One animal per panel. See Fig. 6a in main text.

Supplemental Material S19: Importance and spatial information; Related to Fig. 5

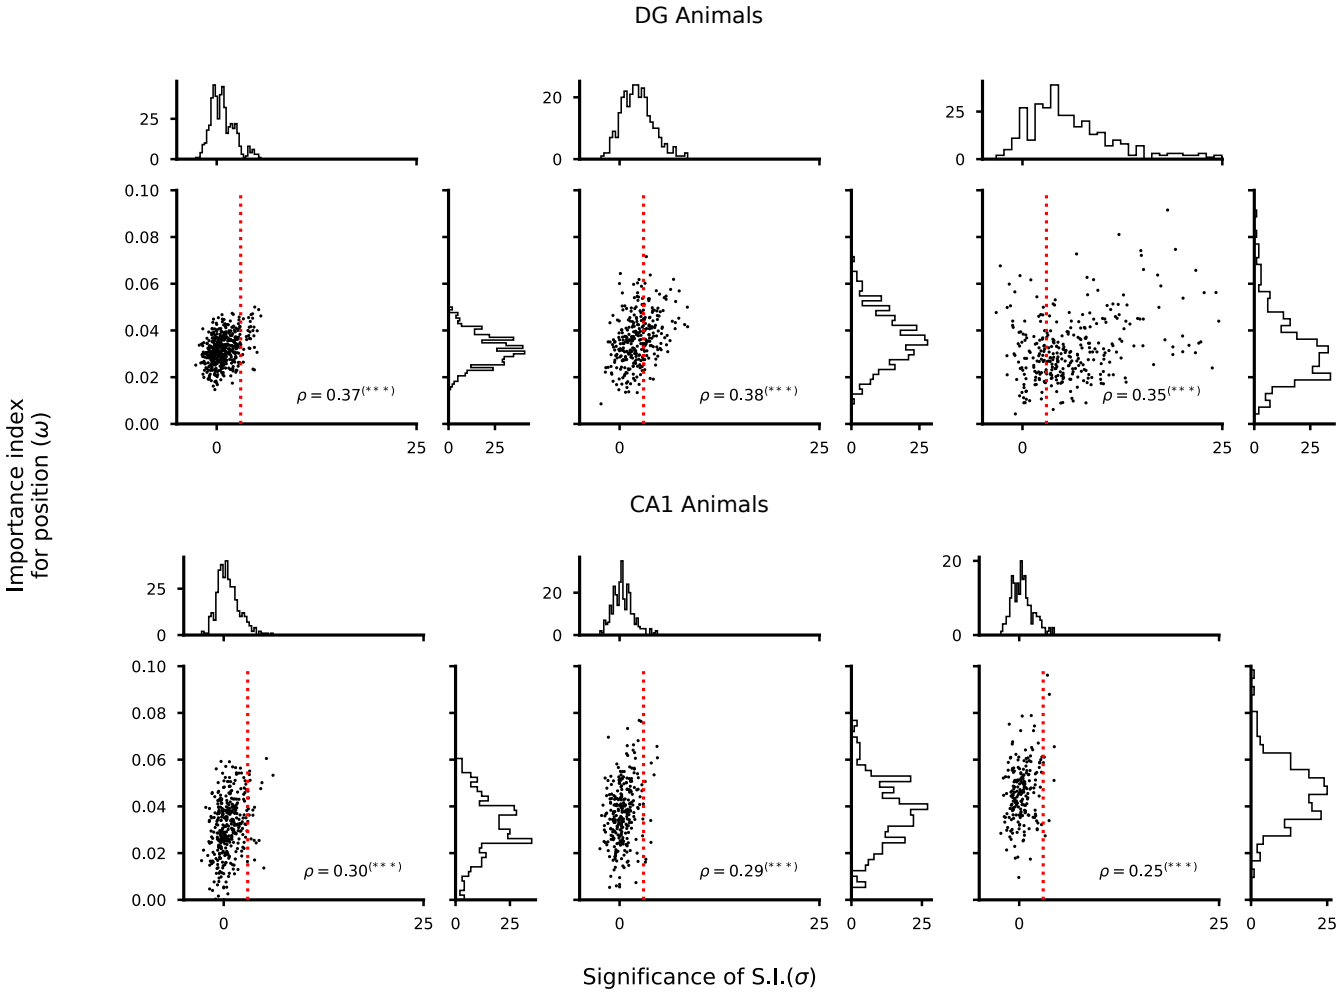

Figure S19: Importance and spatial information. Relation between importance index and significance of spatial information. Top: DG animals. Bottom: CA1. One animal per panel. See Fig. 5 in main text.

Supplemental Material S20: Importance for direction and direction information; Related to Fig. 6

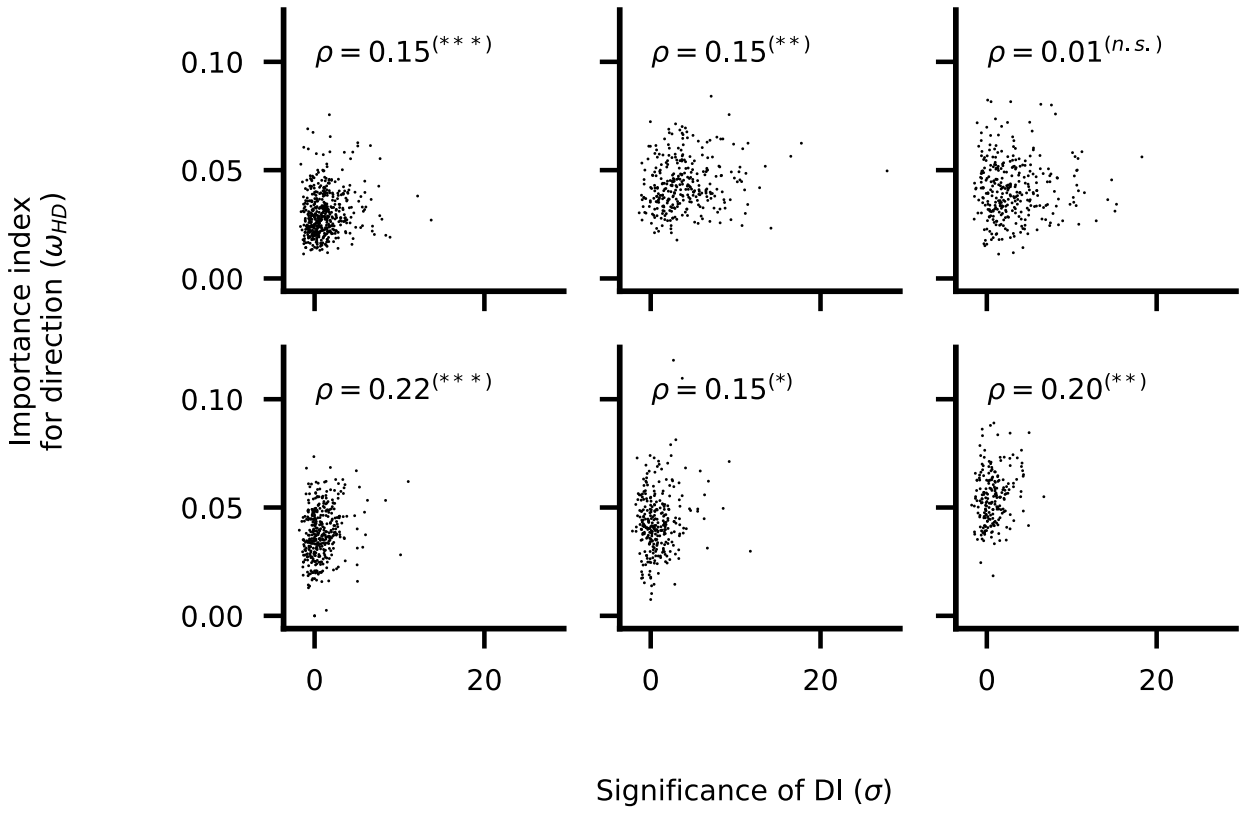

Figure S20: Importance for direction and direction information. Relation between importance index for direction of motion and significance of direction information. Top: DG animals. Bottom: CA1. One animal per panel. See Fig. 6 in main text.

Supplemental Material S21: Importance for position and importance for direction; Related to Fig. 7

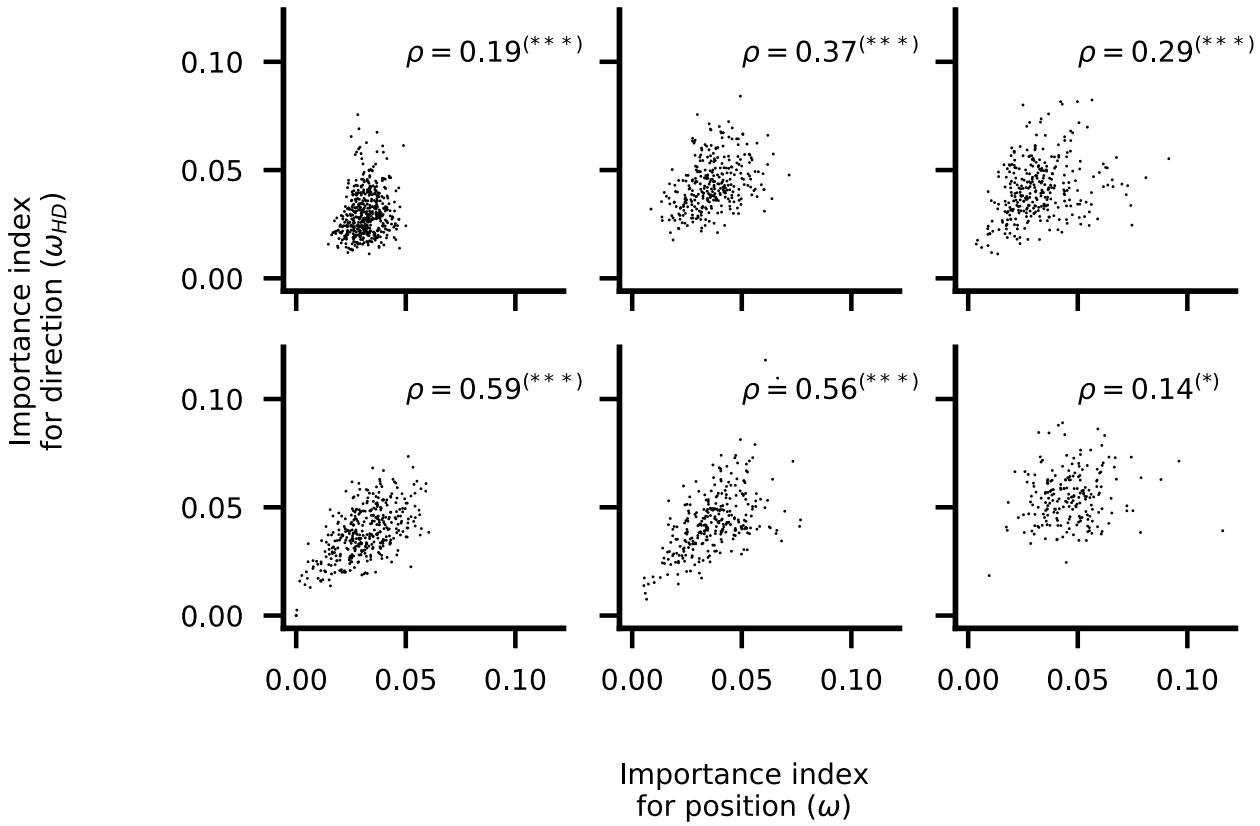

Figure S21: Importance for position and importance for direction. Relation between importance index for position and importance index for head direction. Top: DG animals. Bottom: CA1. One animal per panel. See Fig. 7 in main text.

Supplemental Material S22: Histology and lens placement

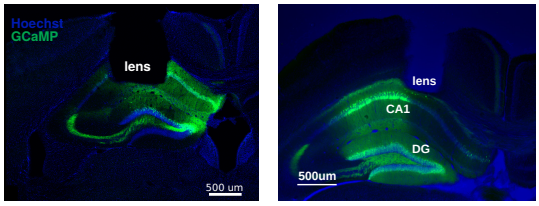

Figure S22: Histology and lens placement. Confocal image for DG (left) and CA1 (right). For all the mice presented in this report the histology confirmed the adequate placement of the lens.

Supplemental Material S23: Relation between correlations and importance index; Related to Fig.4,8

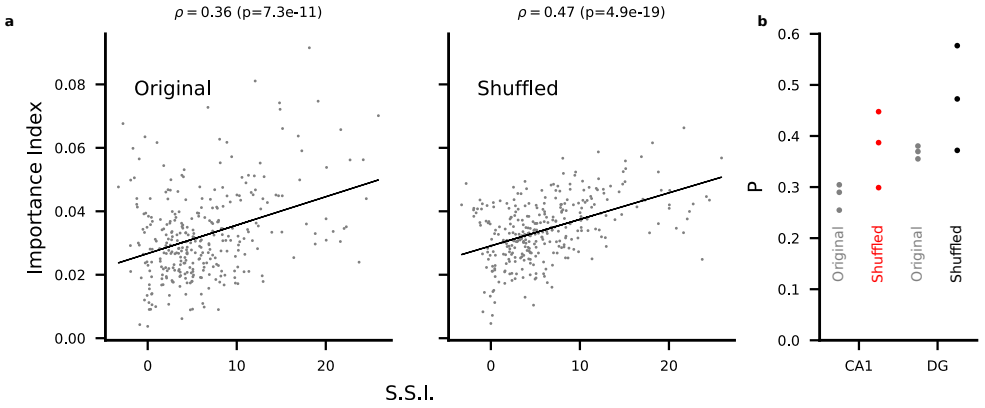

Figure S23: Relation between correlations and importance index. To study the relation between correlation and importance index, we looked at the importance index of single cells before and after destroying correlations among cells. In this way we could measure the independent contribution of single cells to the encoding of position since correlations were destroyed. On the one hand, after our shuffling procedure we would expect that the contribution of a cell to position encoding should be more tied to single cell properties such as spatial tuning. On the other hand, the shuffling procedure does not destroy the correlations that are induced by the relative tuning of cells to a given spatial location, often referred to as signal correlations, and therefore we would still expect some relation between importance index and spatial tuning. For instance, if multiple cells encode the same location, we would expect their importance index to be generally lower than that of cells tuned to location where no other cells are tuned to. Indeed, in our data we found that in most of our animals the correlation between importance index and significance of the spatial information becomes stronger after destroying correlations, suggesting that single cells properties become more relevant in the case of independent coding. However, the effect is marginal so that also in after destroying correlations both place and non-place cells contribute to accurate position decoding. Therefore, also after destroying correlations single cell properties don't fully explain importance index. Finally, we'd like to stress that in Fig. 4c we decoded position from non-place cells alone, defined as cells that do not pass a threshold for location-specific tuning. These results show that our findings are not an artifact of the choice of decoding position but are rather a manifestation of the distributed nature of the spatial code. **a**) Scatter plot of importance index in relation to the significance of spatial information (SSI) for all cells of a representative animal DG3. Left: original data. Right: same as in the left panel but with data in which correlations were destroyed through shuffling. P-values for Pearson-r correlations between the two quantities are reported on top of each panel. **b**) Values of Pearson-r correlations between importance index and SSI in original data (grey dots) and shuffled data (red, CA1; black, DG) for all the 6 mice analyzed. For the shuffled data, mean values of importance index across 20 shufflings was used.

Supplemental Material S24: Control for poorly segmented cells; Related to Fig.4,8

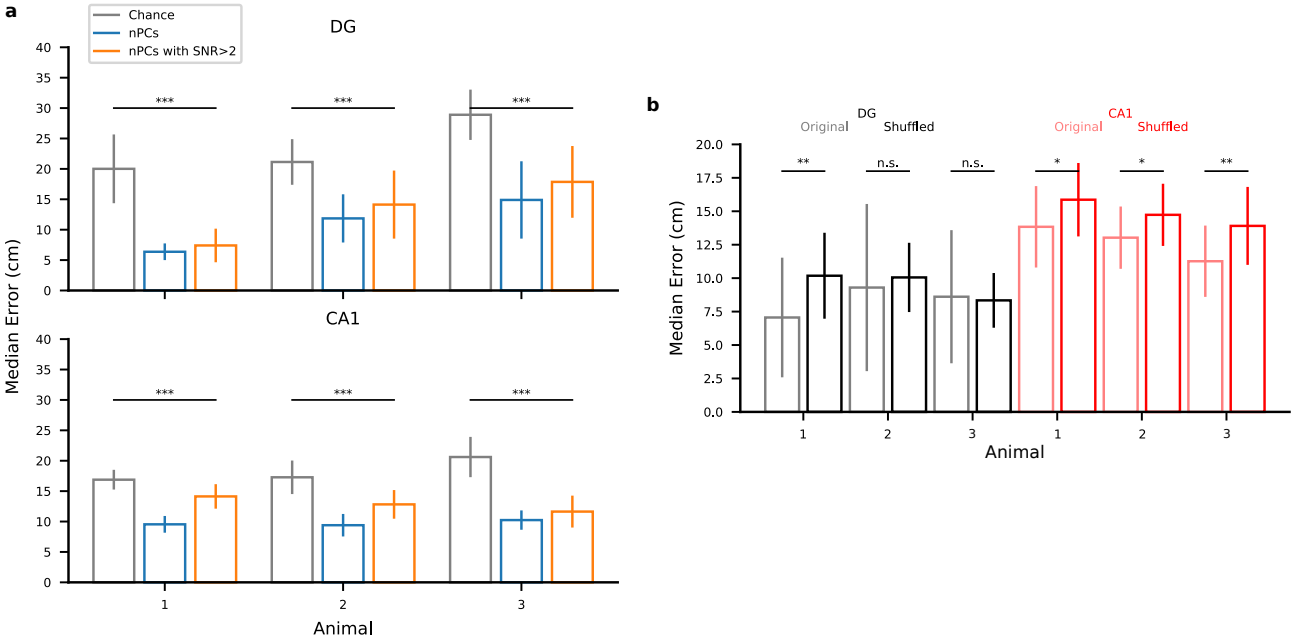

Figure S24: Control for poorly segmented cells. One may wonder if the main conclusions are due to the fact that most of the weakly spatially tuned cells may result from poor unit discrimination in the FOV. We rejected this hypothesis be removing the poorly discriminated cells and verifying that position could be decoded from non-place cells and that destroying correlation has still an effect in CA1.

To isolate the cells resulting from poor discrimination we used the signal-to-noise-ratio between the raw florescence and the deconvolved calcium trace:

$$SNR = \frac{||C||^2}{||C - C_{raw}||^2}$$

where C is the denoised calcium trace and C_{raw} is the raw calcium signal after background removal. This quantity has been introduced by the authors of the CNMF-e algorithm to assess the quality of the extracted signals (Zhou et al., 2018). We then looked at the relation between the place code and the quality of the segmentation. To verify that our main conclusions were not due to poorly identified cells classified as non-place cells, we introduced a criterion of a minimal signal to noise ratio of SNR=2. Our main results that position could be decoded from non-place cells and that destroying correlations had an effect on CA1 remain valid when we selected only neurons with a signal to noise ratio larger than 2. **a**) Position could be decoded from non-place cells also when the analysis is restricted to non-place cells whose SNR is higher than 2. Here we show the decoding position error for all DG (top) and CA1 (bottom) animals for non-place cells (blue) and also for non-place cells for which the signal-to-noise ratio (SNR) after source extraction from the video was higher than 2 (Mann-Whitney U test for significance, ***p<0.001). **b**) Decoding position error for original (light colors) and shuffled data (bright colors) after destroying correlations. Only cells with an SNR>2 were used in the decoding (28%, 91%, 98% in the 3 DG mice; 21%, 22%, 82% in the 3 CA1 mice). We therefore conclude that although untuned cells in our data are overall more likely to have a low SNR as result of poorer signal extraction than tuned cells, the main conclusions are not due to this type of artifacts in the imaging technique.

Table S1

	Authors	Notes	Mean rate
CA1	(Frank et al., 2001)	R, Ep, T-maze	0-5Hz
	(Lee et al., 2004)	R, Ep, enriched circular track	mean 1.1Hz, max 8.7Hz
	(Kelemen and Fenton, 2010)	R, Ep, shock av.	0-2Hz
	(Mizuseki et al., 2012)	R, Ep, multiple taks	(0.88±1.23)Hz
	(Bittner et al., 2017)	R, Ep, 1D	25Hz (peak in field)
	(Ziv et al., 2013)	M, Im, 1D	0.1-0.5 ev/s
	(Xia et al., 2017)	M, Im, two-chamber	∼0-0.2 ev/s
	(Allegra et al., 2019)	M, Im (2p), 1D, virtual,	0.1-0.05 ev/s
	Our work	M, Im, open field	0.06 ± 0.04 ev/s
DG	(Jung et al., 2093)	R, Ep, 8-arm	0-0.2Hz
	(Nitz and McNaughton, 2004)	R, Ep, open field,	15-30Hz (10-20Hz in CA1)
	(Leutgeb et al., 2007)	R, Ep, open field, lots of mossy	0-10 Hz (mean in-field) 5-25 Hz (peak)
	(Neunuebel and Knierim, 2012)	R, Ep, square	5-20 Hz (peak)
	(Pilz et al., 2016)	M, Im (2p), wheel,	0.01 ev/s
	(Danielson et al., 2016)	M, Im (2p), wheel, threshold-based inference	0-0.003 ev/s
	(Allegra et al., 2019)	M, Im (2p), 1D, virtual, higher selectivity in DG	0.1-0.05 ev/s
	Our work	M, Im, open field	0.15 ± 0.12 ev/s

Table S1: Firing rates in CA1 and DG. Related to Fig.4 Mean rates on CA1 and DG cells in spatial tasks (R=rats, M=mice, Ep=e-phys, Im=calcium imaging). One of the observations of our study is that the average rate of events in DG is larger than in CA1. Here we review the literature to show that our observation is compatible with the values that are reported, especially if one considers the high variability across different studies. Differences in behavioral protocols may strongly affect overall activity. This is especially true for a region like the hippocampus that is highly involved in a host of cognitive processes integrating information from virtually all the sensory areas and having a big role in memory. Most of the studies on the hippocampus in spatial tasks focus on 1-dimensional (1D) tracks (wheels or treadmills) where rats or mice are trained to collect rewards at the ends of a treadmill (e.g., experiments from the Buzsaki, Ziv, Schnitzer, Golshani, Dombeck, Losonczy labs among others) or specific mazes with several explicit cues (Knierim, Frank labs). These protocols are specifically designed to study particular properties of the activity of hippocampal cells and it may be inadequate to directly compare them to our study. Studies in freely moving rats include shock-zone avoidance (Kelemen and Fenton, 2010) and tasks combining the location of stimuli with other properties of the task such as identity and memory of previous trials (Wood et al., 1999). A relatively smaller number of studies has considered rats foraging in open fields (e.g., experiments from the Leutgeb, Moser, McNaughton labs) with a relatively small number of explicit or controlled cues. The reported quantities are also defined in various ways, for instance firing rates may be reported as mean rates across sessions (often of the order of 1 Hz) or as peaks within field (often of the order of 5-20Hz – in addition while fields are defined more or less in similar ways, cells with

significant fields are defined in inconsistent ways across studies). Furthermore, they may refer to different behavioural states (e.g., locomotion or immobility) or averaged across states. Thus, the firing properties of CA1 and dentate gyrus neurons can vary widely depending on the parameters of the experimental protocols. Differences in recording techniques are also important and they are probably one of the main sources of variability in estimates of neural activity. The rate of events reported in calcium imaging studies is always significantly smaller than the rate of spikes recorded using electrophysiological techniques. This is not surprising given that the techniques are very different (calcium is only a proxy of neuronal spiking) and often the neurons recorded using calcium imaging are of different types (see for example (Harvey et al., 2009; Dombeck et al., 2010; Ziv et al., 2013)). In electrophysiology, cells are isolated through clustering and selected using properties such as rates and shape of the action potential. A relatively large number of spikes is needed to isolate cells, introducing a bias toward most active neurons. Some labs include recordings during sleep to increase their statistical power but regions such as the dentate gyrus may express different activity profiles during sleep (Neunuebel and Knierim, 2012). In calcium imaging, cells are selected by viral expression and morphology. This allows to visualize the isolated cells, but spiking activity can only be inferred from the calcium signals. Through the relevant literature, we observe that electrophysiology studies typically report firing rates that are one order of magnitude higher than "calcium event rates" in imaging studies. Due to the relatively low sampling rate (5 Hz in our study, up to 30 Hz in other studies) and the slow temporal profile of the florescence profiles, spikes may go undetected in calcium imaging. Hence, it is highly likely that calcium events are in fact due to multiple spikes collected in a short time frame. Also, several spikes may simply be lost throughout the session due to low signal-to-noise. Lower event rates for calcium imaging with respect to electrophysiology are therefore expected. Important differences exist also among calcium imaging techniques, more specifically between one (or single) and two-photon imaging. One crucial step in the processing of calcium imaging videos is that of segmentation, i.e., the identification of the regions of interest (ROIs) in the field of view from which to compute the fluorescence signals and hence each putative cell's activity. Cell segmentation is very different in one and two-photon imaging. One-photon imaging relies on algorithms such as CNMF-e, which are sensitive to the statistics of cellular activities. Higher activity generally leads to improved source separation because of the resulting higher signal to noise ratio (SNR) in fluorescence signal, hence one-photon imaging may have a bias towards more active cells. Instead, in two-photon imaging cells are more easily identified, even in the case they are silent, and so manual segmentation or automated techniques only based on spatial profiles (e.g., apparent shape of the ROI) can be used. The bias of one-photon imaging is partially reduced by the fact that CNMF-e uses both spatial and temporal dynamics profiles to separate signal sources, hence typically even cells with low signal to-noise ratio and low activity are identified as separate sources, but a bias may still exist compared to two-photon imaging. The CNMF-e method has quickly become the *de facto* standard in the mini-endoscope imaging field, and now it is also been used in two-photon imaging studies together with other tools for automated segmentation (Pachitariu et al., 2017). Although CNMF-e data is now available for both techniques, it is still difficult to compare one and two-photon imaging because two-photon imaging is used in head-fixed preparations (e.g., mouse on a treadmill) rather than in freely-moving animals as one-photon imaging allows for. Still, it is possible that one-photon techniques underestimate the fraction of inactive cells, with a consequent inflation of the fraction of active cells and place cells (see also (Danielson et al., 2016)). An additional note should be made about methods of inference of cellular activity from fluorescence signals. This operation typically consists of many steps and has been rapidly evolving in the last few years thanks to the introduction of automated methods necessary for large scale imaging (from hundreds to tens of thousands of cells). While earlier studies applied thresholding to the fluorescence signal (typically measured as changes over an estimated baseline, $\Delta F/F$) to estimate cellular activity in time, more recent methods use more sophisticated approaches based

on models of calcium dynamics to get a more precise estimation of cellular activity (Dombeck et al., 2007; Mukamel et al., 2009; Pnevmatikakis et al., 2016; Pachitariu et al., 2017; Zhou et al., 2018). Thresholding may obviously lead to underestimating cellular activity since one calcium event may be caused by multiple spikes closely occurring in time like in bursts. Ziv and colleagues used a different approach whereby a positive trend in fluorescence signal was associated to spiking activity (Ziv et al., 2013). These different preprocessing choices constitute a further element of variability in the reported data, especially since calcium dynamics is not always known with high precision, depending on cell types and calcium indicators used, among other factors. More research will be needed to evaluate calcium imaging techniques against a ground truth of intra- or juxta-cellular recordings. This is a very active field of research and the data collected with these techniques may still reflect their methodological variability. With all these premises in mind, it is anyway useful to review the literature in which the authors report the activity of neurons in the dentate gyrus and in CA1. The firing properties of dentate gyrus granule cells are still relatively poorly understood. Most of the works that report sparsity arise from immediate gene expression studies and from rat studies by the Moser, Leutgeb, McNaughton and Knierim labs (among others). Notice that sparsity is often used with different meanings. Here we refer to the fraction of cells that are active, and not to the average activity of each neuron. Recent work from Goodsmith and Knierim (GoodSmith et al., 2017) has highlighted the difficulty in identifying granule cells using tetrode recordings. Based on cell morphology, histology and viral expression, we are highly confident that our signals correspond to the activity of granule cells of the dentate gyrus. We therefore believe that our study is the first to show dentate gyrus granule cells activity in freely moving foraging task in an open field. This makes any comparison with previous literature more difficult. Among the literature, we highlight a few recent works worth of note. A recent work by Allegra and colleagues (Allegra et al., 2019) finds very similar calcium activity profiles for CA1 and DG cells, though with a much lower spatial selectivity for CA1 cells. One of the most careful assessment of DG GCs activities (GoodSmith et al., 2017) finds rates for DG similar to CA1 with a 20-30% ratio of place cells in DG. Nitz and McNaughton (Nitz and McNaughton, 2004) report higher rates for DG cells while rats explored a novel rather than a familiar environment. This we report all the most relevant results from the literature. Overall, our observations are compatible with existing literature within the variability across studies. Notice that the differences between DG and CA1 in terms of overall activity are much smaller than the variation across studies. We therefore believe that the question of which area is more active in spatial tasks and what the ratios of place cells are in these regions remain open. We believe that imaging studies like ours targeting such comparison may strongly contribute to strengthening our understanding of the different computational properties of CA1 and DG (see for example (Allegra et al., 2019)). In conclusion, although differences between DG and CA1 place cell number may depend on several other factors such exposure times during habituation, our analysis suggests that the main drive for place cell counts and differences between these two areas comes from the different activity levels. In particular, a non-negligible proportion of CA1 cells are more sparsely active than previously reported and therefore CA1 place cells may have been over represented in place cells studies. We encourage researchers to adopt a more systematic assessment of the significance of the spatial information as we suggest in our work and to always report at least the proportion of place and non-place cells that are found.